

---

# Firely Server

**Firely**

**Feb 13, 2023**



# FIRELY PRODUCTS

<b>1</b>	<b>Overview of Firely Server, Plugins and Facades</b>	<b>3</b>
1.1	Framework . . . . .	3
1.2	Complete Server . . . . .	4
1.3	Plugins . . . . .	5
1.4	Facades . . . . .	5
<b>2</b>	<b>Getting Started</b>	<b>7</b>
2.1	Configuration . . . . .	7
2.2	Running Firely Server . . . . .	7
2.3	Access Firely Server . . . . .	8
<b>3</b>	<b>Release notes and upgrade management Firely Server</b>	<b>15</b>
3.1	Current Firely Server release notes (v4.x) . . . . .	15
3.2	Old Firely Server release notes (v3.x) . . . . .	37
3.3	Old Firely Server release notes (v2.x and earlier) . . . . .	54
3.4	Security notifications for Firely Server . . . . .	75
3.5	How to upgrade Firely Server? . . . . .	76
3.6	Firely Server Roadmap . . . . .	80
<b>4</b>	<b>Frequently asked questions</b>	<b>83</b>
4.1	Conflicting resources upon import . . . . .	83
4.2	Searchparameter errors for composite parameters . . . . .	83
4.3	.NET SDK not found . . . . .	84
4.4	Homepage takes long to load . . . . .	84
<b>5</b>	<b>Configuring Firely Server</b>	<b>85</b>
5.1	Firely Server settings . . . . .	85
5.2	Firely Server settings with Environment Variables . . . . .	95
5.3	Configure the Administration API . . . . .	97
5.4	Managing Conformance Resources . . . . .	99
5.5	Validating incoming resources . . . . .	104
5.6	Using the In-Memory storage . . . . .	105
5.7	Using MongoDB . . . . .	106
5.8	Using SQL server . . . . .	110
5.9	Using SQLite . . . . .	114
5.10	Configure http and https . . . . .	117
5.11	Cross Origin Resource Sharing (CORS) . . . . .	118
5.12	Log settings . . . . .	118
<b>6</b>	<b>Firely Server deployment options</b>	<b>125</b>
6.1	Using Firely Server on Docker . . . . .	125

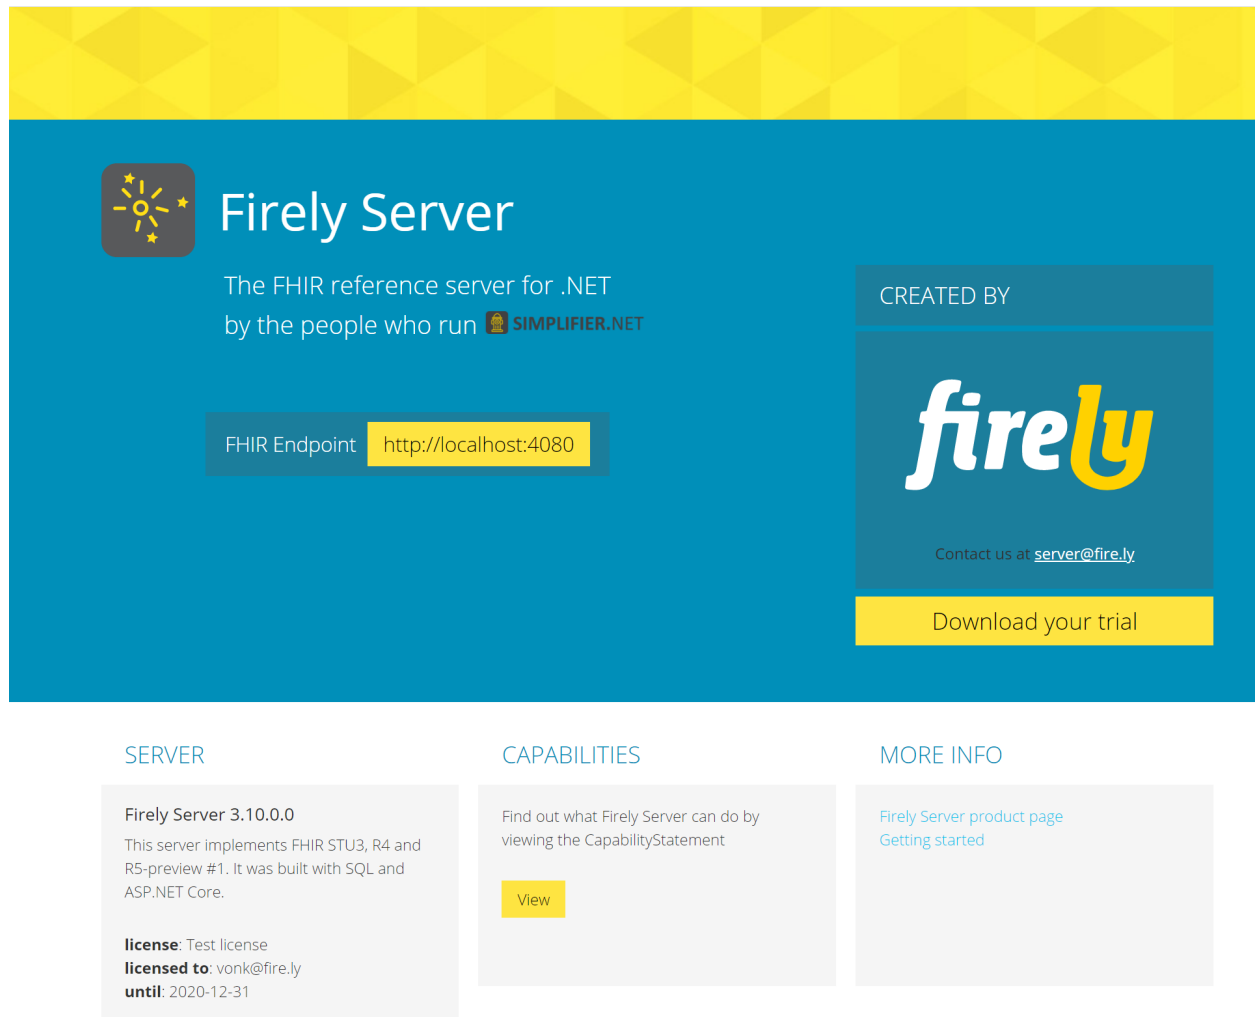
6.2	Using Firely Server on Kubernetes with a Helm chart . . . . .	133
6.3	Firely Server deployment on Azure Web App Service . . . . .	134
6.4	Firely Server for your Simplifier project . . . . .	137
6.5	Deploy Firely Server on a reverse proxy . . . . .	137
6.6	Set up an Identity Provider . . . . .	148
6.7	Performance of Firely Server . . . . .	149
<b>7</b>	<b>Features</b>	<b>157</b>
7.1	FHIR RESTful API . . . . .	157
7.2	Bulk Data Export . . . . .	162
7.3	Patient \$everything . . . . .	166
7.4	\$lastN Observations . . . . .	167
7.5	Custom Operations . . . . .	174
7.6	Custom Resources . . . . .	177
7.7	Terminology services . . . . .	178
7.8	Custom Search Parameters . . . . .	184
7.9	Errata to the specification . . . . .	186
7.10	Subscriptions . . . . .	187
7.11	X-Provenance header . . . . .	189
7.12	Reset the database . . . . .	190
7.13	Preloading resources . . . . .	191
7.14	Multiple versions of FHIR . . . . .	191
7.15	\$erase resources . . . . .	195
7.16	Liveness and readiness . . . . .	196
<b>8</b>	<b>Security</b>	<b>199</b>
8.1	Access control and SMART on FHIR . . . . .	199
8.2	Auditing . . . . .	211
<b>9</b>	<b>Firely Server Administration API</b>	<b>217</b>
9.1	Functions . . . . .	217
9.2	Configuration . . . . .	217
9.3	Database . . . . .	217
<b>10</b>	<b>Compliance with national ImplementationGuides and certification programs</b>	<b>219</b>
10.1	Nictiz . . . . .	219
10.2	GDPR compliance . . . . .	220
10.3	HIPAA compliance . . . . .	221
<b>11</b>	<b>Programming API Reference &amp; Plugins</b>	<b>225</b>
11.1	Firely Server Plugins . . . . .	225
11.2	Firely Server Facade . . . . .	247
11.3	Plugins available for Firely Server . . . . .	267
11.4	Important classes and interfaces . . . . .	291
11.5	Architecture . . . . .	313
11.6	Dependencies of Firely Server and their licenses . . . . .	314
<b>12</b>	<b>Tools</b>	<b>317</b>
12.1	Firely Server Ingest (FSI) . . . . .	317
12.2	Firely Auth . . . . .	323
<b>13</b>	<b>OpenAPI</b>	<b>349</b>
13.1	Limitations of Capability Statement to OpenAPI mapping . . . . .	349
13.2	Generating OpenAPI for a local instance of Firely Server . . . . .	349

<b>14 Contact us</b>	<b>351</b>
<b>Index</b>	<b>353</b>




Firely Server is Firely's FHIR server. It was formerly named 'Vonk' and is the successor to our Spark server.

Firely Server is the answer to the growing need for a stable server that can be used in a variety of production environments. A public sandbox is available at <https://server.fire.ly>, which is free to use and intended for testing and educational purposes only. For convenience, the public server API can be explored using our Swagger-based [web UI](#).



The landing page for Firely Server features a yellow header and a blue main section. On the left, there is a logo of a sun with rays and the text 'Firely Server'. Below the logo, it says 'The FHIR reference server for .NET by the people who run SIMPLIFIER.NET'. A box displays the 'FHIR Endpoint' as 'http://localhost:4080'. On the right, a box says 'CREATED BY' followed by the 'firely' logo. Below the logo, it says 'Contact us at [server@fire.ly](mailto:server@fire.ly)' and a yellow button labeled 'Download your trial'. At the bottom, there are three columns: 'SERVER' with version '3.10.0.0' and license information, 'CAPABILITIES' with a 'View' button, and 'MORE INFO' with links to the product page and getting started.

**Firely Server**

The FHIR reference server for .NET  
by the people who run  SIMPLIFIER.NET

FHIR Endpoint `http://localhost:4080`

CREATED BY

**firely**

Contact us at [server@fire.ly](mailto:server@fire.ly)

Download your trial

**SERVER**

Firely Server 3.10.0.0  
This server implements FHIR STU3, R4 and R5-preview #1. It was built with SQL and ASP.NET Core.

**license:** Test license  
**licensed to:** [vonk@fire.ly](mailto:vonk@fire.ly)  
**until:** 2020-12-31

**CAPABILITIES**

Find out what Firely Server can do by viewing the CapabilityStatement

[View](#)

**MORE INFO**

[Firely Server product page](#)  
[Getting started](#)

On these pages we provide you with the documentation you need to get up and running with your own standard Firely Server installation, as well as information on how to contact us when you have additional needs, such as a custom implementation or support contract.

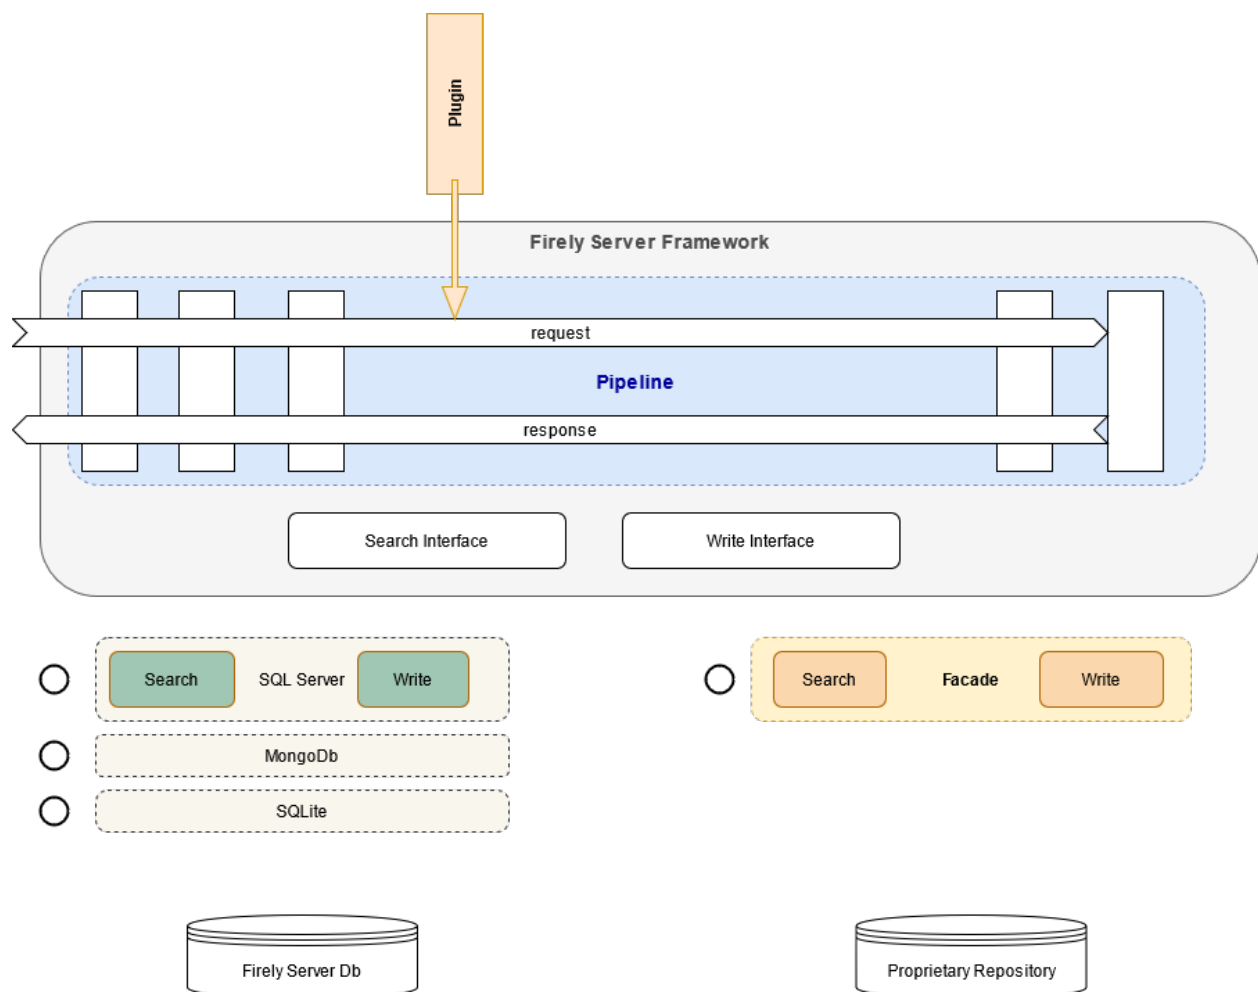




## OVERVIEW OF FIRELY SERVER, PLUGINS AND FACADES

### 1.1 Framework

Firely Server is not just a FHIR Server, it is a processing pipeline for handling standard and custom FHIR requests. *Firely Server* consists of this pipeline filled with processors to handle the interactions defined in the FHIR RESTful API. With *Plugins* you can add your own processors to the framework to perform custom operations, or fill in cross-cutting concerns for your business. A *Facade* is a type of plugin that provides a data access layer for an existing data repository (e.g. a proprietary system). This image sums it all up:



Firely Server comes in several editions:

- Sandbox: try it right now on <https://server.fire.ly>
- Evaluation: get an evaluation license from [Simplifier.net](https://simplifier.net), allowing you to explore all functionality for free during a week (renewable)
- Community: use Firely Server for free, but only on SQLite
- Commercial use with professional support in different tiers: Startup or Scale

For more information and pricing visit the [product site](#).

## 1.2 Complete Server

Firely Server is a FHIR Server out of the box. It is built with Microsoft .NET Core and runs on any of the platforms for which a [.NET Core Runtime](#) is available. Linux, Windows, MacOS, Docker etcetera. Installation can be done in minutes. After that you can configure main features and further details:

- Choose your *database*: [SQLite](#) is configured by default, but for serious use you'd want to configure [MongoDB](#) or [SQL Server](#).
- Configure the level of *validation*: Firely Server can be very loose or very strict on the validity of the resources that you send to it.
- Configure *endpoints* for FHIR versions that you want to support (since Firely Server (Vonk) 3.0.0: FHIR STU3 and FHIR R4)
- Fill in your *licensefile*.
- Adjust the *processing* pipeline by trimming it down (excluding certain plugins) or extending it with extra plugins.

Besides configuration of the settings, Firely Server features an [Administration API](#) that allows you to configure the so-called [Conformance Resources](#) that drive parsing, serialization, validation and terminology. The Administration API is pre-filled with Conformance Resources such as the StructureDefinitions, Searchparameters, CodeSystems and ValueSets that come with the FHIR Specification. Beyond that you can use the Administration API to make Firely Server aware of:

- Custom profiles, e.g. national or institutional restrictions on the standard FHIR resources.
- *Custom resources*: you can even define resources beyond those in FHIR and they are treated as if they were standard FHIR resources.
- CodeSystem and ValueSet resources for *terminology*.
- *Custom Searchparameters*: have Firely Server index and search resources on properties that are not searchable with the searchparameters from the FHIR Specification itself.

Read more on Firely Server:

- [Getting Started](#)
- [Features](#)
- [Firely Server deployment options](#)
- [Configuring Firely Server](#)
- [Release notes and upgrade management Firely Server](#)
- [Firely Server Administration API](#)

## 1.3 Plugins

A plugin is a library of code that you can buy, clone or create yourself that implements additional or replacement functionality in Firely Server. Examples are:

- Implementation of a custom operation. E.g. \$document (generate a document Bundle based on a Composition resource), which is available on [GitHub](#).
- Implementation of a cross-cutting concern. Imagine that in your organization every resource that is created or updated must be logged to a very specific location. You may create a plugin that does exactly that.
- Special handling of specific requests. E.g. requests for a Binary resource where you need to merge in binary data from one of your systems.
- Provide custom authentication and authorization methods for compliancy with business or governmental rules.

In all cases, a Plugin is technically a .NET Core assembly (or a set of them) containing well-defined configuration methods that allow Firely Server to:

- add services
- add a processor to the request processing pipeline

Most plugins do both, registering (testable) services that do the actual work and a thin layer around it that adds it as a processor to the pipeline.

Read more on *Firely Server Plugins*.

View the [session on Plugins](#) from DevDays 2018.

## 1.4 Facades

A Facade is a Firely Server processing pipeline working on an existing data repository. That repository could be the database of proprietary system, some API of an existing system or a whole Clinical Data Repository specifically created to open up data through a FHIR API.

The implementation of a Facade is a special type of plugin that registers services to access the existing data repository. By building the data access layer you leverage all of the FHIR processing in Firely Server, connected to your repository - thus creating a FHIR RESTful API for that repository with the least amount of work.

So a Facade is still a Plugin, and therefore technically a .NET Core assembly (or a set of them) having the same well-defined configuration methods. In the case of a Facade it usually only registers services (and no processor), specifically implementing the interfaces that define the data access layer in Firely Server:

- ISearchRepository, for reading and searching
- IResourceChangeRepository: for creating, updating, and deleting

Read more on *Firely Server Facade*.

View the [session on Facade](#) from DevDays 2018.



## GETTING STARTED

If you want to start using the standard Firely Server in your own Windows environment, follow the steps on this page to install and run the server. For non Windows systems, or if you want to use Docker for Windows, please look at the *Using Firely Server on Docker* section.

1. Download the Firely Server binaries and the license file from [Simplifier.net](https://simplifier.net).
2. Extract the downloaded files to a location on your system, for example: C:\FirelyServer. We will call this the working directory.
3. Put the license file in the working directory.
4. In the working directory create a new JSON file and name it `appsettings.json`. You will use this file for settings that you want to differ from the defaults in `appsettings.default.json`. For more background on how the settings are processed, see *Firely Server settings*
5. Open `appsettings.json`, copy the `LicenseFile` setting from `appsettings.default.json` to it and change this property to the name of your license file. For example

```
"License": {  
  "LicenseFile": "firelyserver-trial-license.json"  
}
```

### 2.1 Configuration

The section *Configuring Firely Server* explains how you can further configure the Firely Server.

### 2.2 Running Firely Server

---

**Important:** The next step assumes you have a .NET Core environment installed. If not, please [download and install ASP.NET Core Runtime 6.x.xx Hosting Bundle](#) before you continue. Choose the latest security patch to mitigate security issues in previous versions.

---


When you have completed your configuration changes, start the server. Open a command prompt or Powershell, navigate to your working directory and execute:

```
> dotnet .\Firely.Server.dll
```

### 2.3 Access Firely Server

Firely Server will by default run on port 4080 of the system. Check if Firely Server is running correctly, open a browser and navigate to `localhost:4080`.

You will see the following homepage:



The following links provide documentation and samples for testing the Firely Server restful endpoint:

### 2.3.1 Postman Setup


This article will walk through setting up Postman to access Firely Server restful endpoint.

#### Prerequisites

Sign up for and install Postman. For more information about Postman, see [Get Started with Postman](#).

#### Using Postman



1. Click the following “Fork postman collection into your workspace” link:
2. Click “Fork Collection”



3. Sign-In with your Postman account and click “Fork Collection”. Change the label and workspace names as desired.



4. Variables have been predefined at the collection level for ease of use. Adjust the variables to reflect your Firely Server endpoint and settings and save.

FirelyServer - Restful Examples  FirelyServer - Restful... 

Authorization ● Pre-request Script Tests Variables ● Runs

These variables are specific to this collection and its requests. Learn more about [collection variables](#) ➤

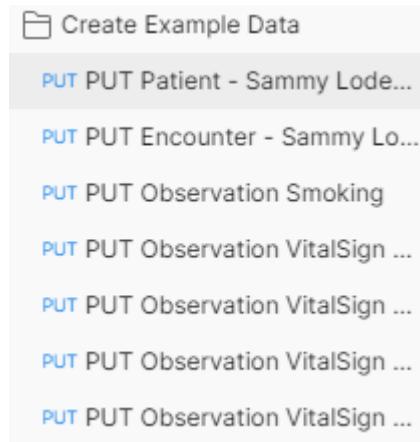
	VARIABLE	INITIAL VALUE ⓘ	CURRENT VALUE ⓘ
<input checked="" type="checkbox"/>	BASE_URL	https://server.fire.ly	https://server.fire.ly

- Test the first request, metadata, as seen below by clicking “Send”. This will return the server Capability Statement.

images/postman\_tutorial\_metadata.png

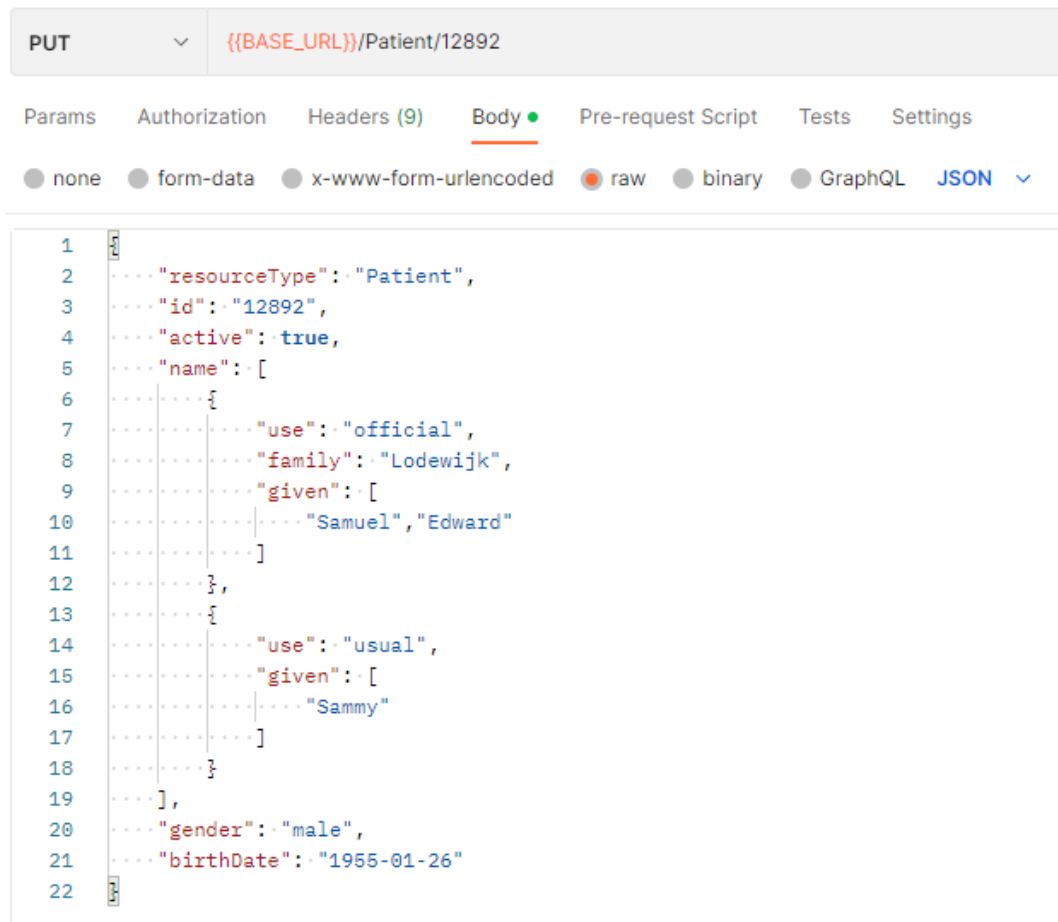
```
{
  "resourceType": "CapabilityStatement",
  "id": "e8bfb522-022d-47d4-a1e6-c05b7fc1175d",
  "meta": {
    "versionId": "de3e4181-97be-4f58-80d2-e01cb2198ef2",
    "lastUpdated": "2022-12-07T19:14:07.7140139+00:00"
  },
  "language": "en-US",
  "url": "http://server.fire.ly/fhir/CapabilityStatement/FirelyServer",
  "version": "1.0",
  "name": "Firely Server 4.10.0 CapabilityStatement",
  "status": "active",
  "experimental": true,
  "date": "2022-12-07T19:14:07.7140229+00:00",
  "publisher": "Firely"
}
```

- The collection folder “Create Example Data” should be executed next. This will setup FHIR records for the basic examples in “Retrieve Example Data” as well as other examples throughout the collection.





7. In the first example PUT `Patient` observe the body tab to see the resource to be created.



8. Continue testing other requests in the collection. The `Features` folder in the collection aligns to the *Features* documentation and provides corresponding live examples.

### 2.3.2 US Core / 21th Century Cures Act

- Firely Server provides full *profile and interaction* support as defined in *Conforming to US Core*:
  - Firely Server can be populated with resources conforming to US Core
  - All elements defined as must-support by the implementation guide are supported
  - All references between FHIR resources defined as must-support by the implementation guide are supported
  - All search and CRUD interactions defined by US Core are supported, including optional search parameters
- Firely provides official support for the following versions of the US Core ImplementationGuide:

US Core Version	Status	References
US Core 3.1.1		<a href="http://hl7.org/fhir/us/core/STU3.1.1/">http://hl7.org/fhir/us/core/STU3.1.1/</a>
US Core 4.0.0		<a href="http://hl7.org/fhir/us/core/STU4/">http://hl7.org/fhir/us/core/STU4/</a>

- All StructureDefinitions for profiles and extensions (v3.1.1) are loaded by default in the standard SQLite administration database of Firely Server. No additional configuration needed in order to validate against these conformance resources.

See [How to Test Firely Server on Inferno](#) for more information on how to pass the official 21st Century Cures Act test.

### Known Limitations

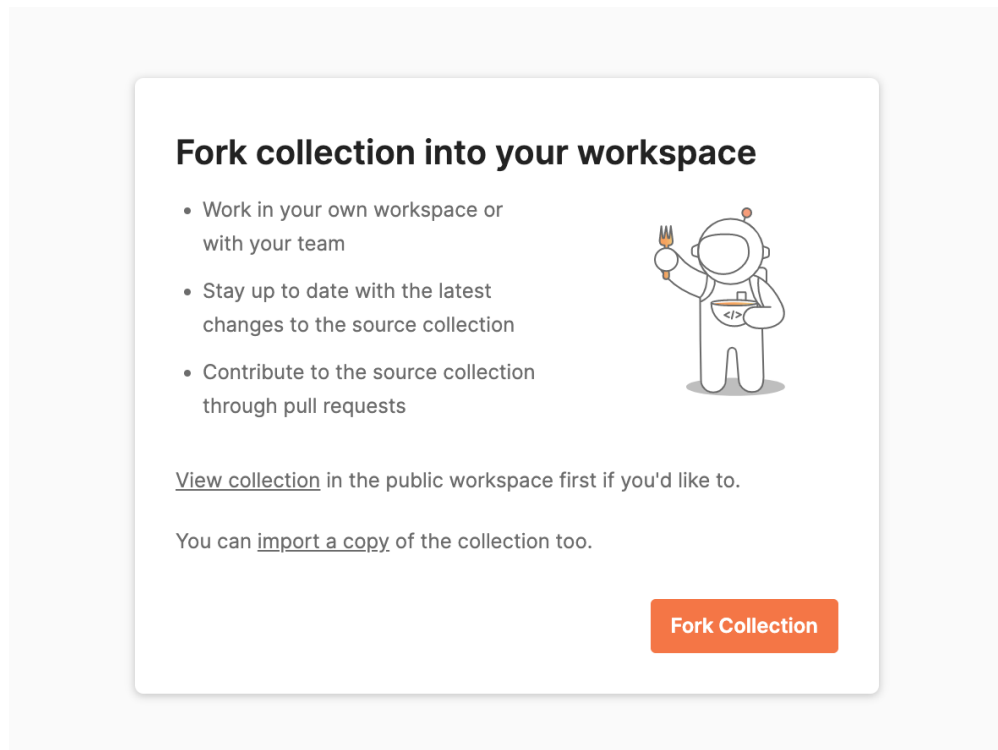
- In order to validate resources claiming to conform to US Core, it is necessary to configure Firely Server to use an external terminology server incl. support for expanding SNOMED CT and LOINC ValueSets. See [Terminology services](#).
- The *on-demand* parameter is not supported for the `$docref` operation on DocumentReference resources

### Test Data

Firely provides test data covering all US-Core profiles and all elements marked as Must-Support. In order to load all examples, two transaction bundles need to be posted against the base endpoint of Firely Server. The following Postman collection provides you with the bundles itself, and the bundle entries as individual PUT requests.

The following steps are necessary in order to execute the test collection against our own Firely Server instance:

1. Select “Fork Collection” or “View collection” in the Postman dialog



2. Sign-In with your Postman account
3. Create a new [Postman environment](#) with a “BASE\_URL” variable and adjust the URL to your server endpoint

Firely Server Public				+ ...
Firely Server Public				Fork
	VARIABLE	TYPE ①	INITIAL VALUE ①	CURRENT VALUE ①
<input checked="" type="checkbox"/>	BASE_URL	default	https://server.fire.ly/	https://server.fire.ly/
	Add a new variable			

4. Make sure that the newly created environment is selected as the active environment
5. Open the collection “Firely Server - US Core Tests”

The screenshot shows the Firely Server interface. On the left is a sidebar with navigation options: Collections (selected), APIs, Environments, Mock Servers, Monitors, Flows, and History. The main area displays the 'Firely Server - US Core Tests' collection, which is expanded to show a list of test requests. At the top right of the main area are 'New' and 'Import' buttons. The list of tests includes:

- POST ONC Example Transaction Bundle (1)
- POST ONC Example Transaction Bundle (2)
- PUT Create US-Core Patient
- PUT Create Group
- PUT Create US-Core Provenance
- PUT Create US-Core Practitioner
- PUT Create US-Core AllergyIntolerance
- PUT Create US-Core CarePlan
- PUT Create US-Core Location
- PUT Create US-Core CareTeam
- PUT Create US-Core CareTeam (2)
- PUT Create US-Core Encounter
- PUT Create US-Core Organization

6. Execute the transaction request, the expected response is “HTTP 200 - OK”.



## RELEASE NOTES AND UPGRADE MANAGEMENT FIRELY SERVER

The following pages contain the release notes of the current major version of Firely Server (v4), as well as, documentation on how to perform an update of the server. See below of older release notes of Firely Server.

We recommend that you familiarize yourself with the documentation on how to upgrade to a new version of Firely Server [How to upgrade Firely Server?](#) before proceeding. Please take a look at our [Firely Server Roadmap](#) to see which features we are planning to work on next.

### 3.1 Current Firely Server release notes (v4.x)

#### 3.1.1 Release 4.10.0, October 6th, 2022

##### Feature

1. HealthCheck: We introduced endpoints for healthchecks, both liveness and readiness, see [Liveness and readiness](#).
2. Provenance: We added support for the X-Provenance header, see [X-Provenance header](#).

---

**Note:** If you use this release as a facade implementation, you need to disable the provenance header configuration in the pipeline.

---

3. Compartments: The reference searchparameter `Binary.securityContext` can now be used as a link for the Patient compartment. If SMART on FHIR is enabled and a “patient”-level scope is being used, Binary resources must now point to a Patient resource matching the “patient” claim in the access token. For facade and plugin implementations this change means that the `ICompartment` information is now available for requests on the `Binary` resource type.
4. Audit on search: We enhanced audit logging by exposing the literal references of the current page of the results from a search request. For file logging, this information can be included by adding `SearchResultSummary` to the [output template](#). In `AuditEvent` resources, each search result (reference) is included as an `entity` node, in addition to the entity node for the bundle with the query parameters.
5. Search: We added support of `:of-type` search modifier for search parameters targeting Identifier elements (e.g. `Patient.identifier`). See the [FHIR spec article on search](#) for more information.

---

**Note:** After the update, all new resources containing Identifier elements will be automatically indexed and become searchable with the `:of-type` modifier. If you want to make discoverable the data you already have, you would need to [re-index](#) your data for some of the search parameters you want to use with the `:of-type` modifier.

---

6. Search: Added support for the `Prefer` header when searching, with allowed values `handling=strict` or `handling=lenient`, as per [the specification](#).
7. Search: Added support for wildcards on `_include=*` and `_revinclude=*`, so you can (rev)include all linked resource at once, as specified with [search](#).
8. Search: Canonical references can now be searched using the `:below` modifier, that will match a prefix of the version, as specified for [References and versions](#)
9. US Core: Added support for the `$docref` operation. Firely Server can retrieve all existing `DocumentReferences` for a Patient, but cannot yet generate a document. This operation is defined as part of the [US Core implementation guide](#).
10. Mime-type: If a request did not specify a mime-type in the `Accept` header, Firely Server would default to `application/fhir+json`. If you prefer XML, this default can now be set in the [Response options](#).
11. [Preloading resources](#) is updated and now works for all FHIR versions. Please note that this feature is still meant for limited amounts of (mainly example) data. For loading large amounts of data we recommend to use [Firely Server Ingest \(FSI\)](#).

## Database

1. We introduced new optimizations for the **MongoDB** data schema and queries. These optimizations will improve search performance for elements of type `dateTime` and `decimal`. Please read below notes for the upgrade process.

**Attention:** The upgrade procedure for Firely Server running on MongoDB requires a mandatory migration. If your collection contains a lot of resources, this may take a very long time. Therefore, the MongoDB upgrade script has to be executed manually. The script can be found in *mongodbFS\_SchemaUpgrade\_Data\_v22\_v23.js*

Here are some guidelines:

- We tested it on a MongoDB collection with a size of 500GB. The upgrade script took around 24 hours to complete on a fairly powerful machine.
- As always, make sure you have a backup of your database that has been tried and tested before you begin the upgrade.
- Please make sure that Firely Server is shutdown before you execute the script.
- If you encounter problems running the script, or need any assistance, please [contact us](#).

The update script will update the data that is stored in the database. Although Firely Server can be started as soon as the migration is finished, it will have decreased performance during the first day of operation. This is due to a change in indexes which requires them to be rebuilt in the background.

2. We introduced user defined table types in **SQL Server** for an optimization in [Firely Server Ingest 1.4.0](#). The update is in migration script `FS_SchemaUpgrade_Data_v24_v25.sql` and will be applied automatically when `AutoUpdateDatabase=true` in the settings.

## Fix

1. Administration: The order of loading knowledge and conformance resources has been fixed. We made sure that the definitions stored in the administration database take precedence over the definitions from the `specification.zip` file. Any custom implementations of `IModelContributor` are loaded after the database and before the ZIP file.
2. Search: An erratum to the specification of R4 has been made, changing the type of search parameter `Resource-profile` from `uri` to `reference` (with target `StructureDefinition`). This was an omission in R4 and has been fixed in R5. The change allows searching for `_profile` with the `:above` and `:below` modifier. To take advantage of it, the following steps must be taken:
  - Optionally but recommended: before upgrading, remove the current index data for `Resource._profile` (see [re-indexing](#))
  - Upgrade Firely Server, execute the database migrations and start the server
  - Re-index `Resource._profile` (see [re-indexing](#))

---

**Note:** If you have made manual changes to `SearchParameter/Resource-profile-Fhir4.0` and want to search with the `:above/:below` modifier, you must update your definition to be of type `reference` with target `StructureDefinition`

---

3. Search: Indexing has been fixed for search parameters of type `reference` that index resource elements of type `uri`. The following `SearchParameters` were affected by the bug: Consider [re-indexing](#) your database for these search parameters if you use them.
  - FHIR4: `ConceptMap-source-uri`, `ConceptMap-target-uri`, `PlanDefinition-definition`
  - STU3: `ImplementationGuide-resource`, `Provenance-agent`

---

**Note:** Please note that due to a mistake in the official STU3 specification, search parameters `ConceptMap-source-uri`, `ConceptMap-target-uri` still do not work as expected. The correct search parameter expressions would be `ConceptMap.source.as(uri)` and `ConceptMap.target.as(uri)` while the specification contains `ConceptMap.source.as(Uri)` and `ConceptMap.target.as(Uri)` respectively. The issue has been addressed in R4.

---

4. SMART: With SMART on FHIR enabled, an update-on-create (creating a new resource with an update / PUT) was always denied. This is now fixed.
5. Subscription: if the resthook url in a Subscription did not end with a slash (/), it would get shortened to the last slash in the url. This is now fixed, the whole url is used.

## Plugin and Facade

1. Facade: When building predicates in a Facade implementation of `ISearchRepository` / `IRepoQueryFactory`, exceptions were only translated to the `OperationOutcome`, but not logged. Now they are also logged.
2. API: We will narrow the public programming API in the `Vonk.Core` package in the next major release. To alert you to that, we deprecated the parts that will be removed from the public API.

**Attention:** Please try to build your plugin or facade against `Vonk.Core 4.10.0` to check if you use any of the deprecated parts. If you think some part should not be deprecated, please let us know with a support ticket.

### 3.1.2 Release 4.9.3, September 15th, 2022

#### Fix

1. Starting with Firely Server v4.9.0, a specific search query could fail, with multiple includes on the same parameter, having different type modifier, e.g. *Coverage?\_include=Coverage:payor:Patient&\_include=Coverage:payor:Organization*. That is fixed.

### 3.1.3 Release 4.9.2, August 24th, 2022

#### Fix

1. Starting with Firely Server v4.9.0, validation was only performed against the core specification even if the validation level was set to “Full” and resources sent to Firely Server contained a meta.profile claim.

### 3.1.4 Release 4.9.1, August 1th, 2022

#### Fix

1. Fixed an issue with `_include` and `_revinclude` in case the (rev-)include link was pointing to an element of type “canonical” and not of type “reference”.
2. “\_total” was added as default parameter in the v4.9.0 release. Therefore it must be handled in a facade implementation. The `Vonk.Facade.Relational` package now handles the case of “\_total=accurate”. All other argument values must still be handled in the `ISearchRepository` implementation.
3. Reading the `specification.zip` file from a read-only disk caused an exception.
4. Excluding the `UrlMappingService` from the pipeline configuration and executing a CRUD operation caused an exception.

#### Feature

1. The exposed [SMART capabilities](#) in the `.well-known/smart-configuration` can now be configured in the appsettings. See `SmartAuthorizationOptions.SmartCapabilities` in section [SMART Configuration](#).

### 3.1.5 Public Endpoint Announcement 8 July 2022

The default FHIR version of the [public Firely Server endpoint](#) is now R4.

### 3.1.6 Release 4.9.0, July 6th, 2022

#### Security

1. Upgraded `Microsoft.AspNetCore.Authentication.JwtBearer` dependency as a mitigation for [CVE-2021-34532](#).



## Database

1. Switched the serialization format for decimal types from string to the native decimal type in MongoDB to improve performance.
2. For SQL Server database, if you upgrade Firely Server all the way from v4.2.1, it is likely that the resulting index `vonk.ref.ref_name_relativerelative` differ from a clean installation of Firely Server. The upgrade procedure will try to fix the index automatically. If your database is large, this may take too long and the upgrade process will time out. If that happens you need to run the upgrade script manually. The script for the *admin* database can be found in `sqlserver/FS_SchemaUpgrade_Admin_v22_v23.sql` and the script for the *data* database can be found in `sqlserver/FS_SchemaUpgrade_Data_v23_v24.sql`.

### Attention:

The upgrade procedure for Firely Server running on MongoDB requires a mandatory migration. If your collection contains a lot of resources, this may take a very long time. Therefore, the MongoDB upgrade script has to be executed manually. The script can be found in *mongodbFS\_SchemaUpgrade\_Data\_v21\_v22.js*

Here are some guidelines:

- We tested it on a MongoDB collection with a size of 500GB. The upgrade script took around 24 hours to complete on a fairly powerful machine.
- As always, make sure you have a backup of your database that has been tried and tested before you begin the upgrade.
- Please make sure that Firely Server is shutdown before you execute the script.
- If you encounter problems running the script, or need any assistance, please [contact us](#).

## Fix

1. Fixed an issue where a “/” was missing in the fullUrl of a “search” bundle in case an information model mapping with mode “Path” was used.
2. Fixed an issue where a new resource id was not created when POST was used in a batch or transaction bundle and a resource id was already provided.
3. An invalid system URI was provided by default in `AuditEvent.source.observer.identifier`. Now `http://vonk.fire.ly/fhir/sid/devices|firely-server` is being used to identify Firely Server itself.
4. Adjusted the implementation of conditional create to match the description in <https://jira.hl7.org/browse/FHIR-31965>.
5. Money.currency was not indexed correctly in FHIR R4. Please [contact us](#) if you are using the SearchParameters “price-override” on ChargeItem or “totalgross” / “totalnet” on Invoice. A migration for these fields will be provided upon request. Otherwise, please re-index these SearchParameters. See [Re-indexing for new or changed SearchParameters](#) for more details.
6. Fixed an issue where bundles with conformance claims in meta.profile would have been validated against the profile claims even if the validation level was only set to “Core”.
7. Validating a resource with an element containing only an extension and no value against validation level “Core” will no longer result in an error.
8. SoF: Providing an invalid token to an unsecured operation does not lead to an HTTP 401 error status code. The invalid token is now being ignored.

9. SoF: Fixed unauthorized issue when performing PATCH request with `patient` scope.

### Feature

1. Inferno, the ONC test tool: Firely Server is updated to pass all the tests in the latest ONC test kit (version 2.2.1)! Do you want a demo of this? [Contact us](#).
2. Transactions, including rollbacks, are now fully supported when running Firely Server on MongoDB. Please note that the `SimulateTransaction` setting is no longer available. See [MongoDB Transactions](#) for more details.
3. `$lastN` is now available if Firely Server is running on MongoDB. See [\\$lastN Observations](#) for more details.
4. It is now possible to define exclusion criteria in the appsettings to configure which requests against Firely Server should not be audited. In certain cases, this can reduce the number of captured `AuditEvent` resources. See [Auditing](#) for more details.
5. By default, the `AuditEvent` logging will now include the query parameters sent to Firely Server. These parameters will also be stored in case a request fails (HTTP 4xx or 5xx).
6. The log sinks for `AuditEvent` logging are now configurable in the logsettings. See [Audit log file configuration](#) for more details.
7. Firely Server will throw a startup exception if no default `ITerminologyService` is registered.
8. `CapabilityStatement.rest.resource.conditionalRead` is now set to 'full-support' by default.
9. `_total` is now included in every self-link of a "search" bundle by default.
10. Added support for permanently deleting resources from the database. See [\\$erase resources](#) for more details. You will need an updated license file. Please [contact us](#) if you want to use the feature.
11. Improved the error message in case the JSON serialization format of a FHIR resource does not contain a valid "resourceType" Element.
12. Improved validation in case a non-conformant URI is given in `Quantity.system`. It MUST be a valid absolute URI. In all other cases, a warning will be logged and the element will not be indexed.
13. Improved error message logging in case SQL script fails when the database upgrade is performed automatically by Firely Server.
14. Improved log message in case Firely Server SQL schema needs to be updated by adding the current schema version and the target schema version.
15. Improved access control by no longer allowing retrieval of resources outside of the Patient compartment if SMART on FHIR is enabled and patient-level scopes are provided by the client. Additional resources need to be explicitly allowed by the token.
16. Improved error message in case a condition create/update/delete operation is executed with SMART on FHIR enabled and the client provides a token with limited permissions (e.g. only write-scopes).

### Performance

1. Improved validation performance of large resources. Firely Server will now execute the validation of bundles in a linear amount of time depending on the number of resources in the bundle.
2. Improved performance for chained searches in case SMART on FHIR is enabled.

### 3.1.7 Release 4.8.2, May 10th, 2022

#### Feature

1. A new setting has been introduced in the “Hosting” settings to configure path base. Please check [Firely Server settings page](#) for details.

#### Fix

1. US-Core profiles in conformance resources database *vonkadmin.db* are downgraded from version 4.0.0 to 3.1.1. The upgrade in previous Firely Server was unintentional.
2. CapabilityStatement is cached now based on the absolute request url. With this fix, CapabilityStatement can be properly cached when a request contains **X-Forwarded-\*** headers.
3. For MongoDB repository, set *allowDiskUse* to *true* when using *aggregate* command. This fix solves memory restriction error during aggregation stages (See [MongoDB document](#) for details).

### 3.1.8 Release 4.8.1, Mar 5th, 2022

#### Plugins

1. Upgraded the .NET SDK to 3.8.2. Please review its [release notes](#) for changes.

#### Feature

1. A new option to configure settings regarding TLS client certificates has been introduced in the “Hosting” options. This option allows to set the [ClientCertificateMode](#).
2. Validation of transaction/batch bundles has been enabled by default when posting the resources to the transaction endpoint of Firely Server. Please note that the transaction is executed synchronously. To avoid client timeouts, the default value for the MaxBatchEntries (SizeLimits options) has been reduced to 200.

### 3.1.9 Release 4.8.0, Mar 21st, 2022

#### Plugins

1. Upgraded the .NET SDK to 3.8.0. Please review its [release notes](#) for changes.

#### Database

1. SQL Server
  1. Reduced database size by compressing the resource JSON.

**Attention:** This change requires a complex SQL migration which can be long if you have many resources. To estimate how long it will take for you, you can try running the migration for a subset of your data. The overall migration time will grow linearly with the number of resources in the database.

For our test database containing ~185mln FHIR resources, the migration took approximately 1.5 days.

If you have questions about the migration, please [contact us](#).

### Performance

1. Improved performance for update, `_include/_revinclude` and conditional create interactions

### Feature

1. You can now control the inclusion of the `fhirVersion` mimetype parameter in the Content-Type header of the response. See [Running different versions on different endpoints](#). We chose to change the default for FHIR STU3 to *not* include it as this parameter was introduced with FHIR R4.

### Fix

1. Fixed exception by improving transaction handling when updating and deleting the same resource in parallel.
2. Use correct restful interaction codes in `AuditEvent.subtype` when recording a request to Firely Server
3. `AuditEvent.action` contained the wrong code when recording a SEARCH interaction
4. The name of a custom operation is now recorded in an `AuditEvent`
5. Fixed searching using the `:identifier` modifier in case the identifier system is not a valid URL
6. Searching using a If-None-Exist header was not scoped to an information model, i.e. a request using FHIR R4 also matched STU3 resources
7. Improved error message if `$lastN` operation is enabled but the corresponding repository is not included in the pipeline options
8. Changed `CapabilityStatement.software.name` to Firely Server
9. Fixed SQL Server maintenance job timeouts on large SQL Server databases
10. Improved Bundle reference resolving in some corner cases, which are clarified in the [this HL7 Jira issue](#)

### Security

1. According to the [best practices](#) of docker, Firely Server container runs now under the user and group `firely:firely` instead of running under root privileges.

## 3.1.10 Release 4.7.1, Feb 15th, 2022

### Fix

1. An invalid `CapabilityStatement` was created by Firely Server in case a custom `SearchParameter` overwriting a common `SearchParameter` was loaded, e.g. `“_id”`. `CapabilityStatement.rest.resource.searchParam.definition` contains now the canonical of the more specific `SearchParameter`.
2. The default `CapabilityStatement` contained an invalid canonical in the `.url` element.
3. Enforce referential integrity for the elements `“Composition.patient”` and `“Composition.encounter”` when submitting a document bundle to the base endpoint. The corresponding resources need to be already present on the server (matching based on identifier), otherwise the bundle is rejected.

### 3.1.11 Release 4.7.0, Feb 1st, 2022

**Attention:** With version 4.7.0, Firely Server migrated to .NET 6.0. In order to run the binaries, [ASP.NET Core Runtime 6.x](#) needs to be installed.

#### Feature

1. BulkDataExport is now supported for MongoDB as well. Get started with the [Bulk Data Export documentation](#).
2. Circular references in transaction bundles are now supported. Bundles of type `transaction` and `batch` are permitted to contain resources referencing another resource within the same bundle. This also means that you can now cross reference PUT and POST entries.
3. An option to configure additional token issuers is now available. This is used in settings where the token issuer deviates from the token audience. This new setting replaces the existing `AdditionalEndpointBaseAddresses`. The setting needs to be adjusted manually as it will not be migrated automatically. Please check the [configuration documentation](#) on how to use it.
4. Firely Server now supports receiving document bundles on the base endpoint. Firely Server will extract the narrative of document bundles and store this within a `DocumentReference` resource. Read more about it in the [documentation](#).
5. Added support for transforming *SMART scopes issued by Azure Active Directory* [documentation](#).
6. Firely Server will now recognize the `name` claim in JSON Web Tokens and also include its content in the logs.
7. It is now possible to [provide the Firely Server license via an environment variable](#).

#### Plugins

1. BulkDataExport interfaces were made publicly available in order to provide these to Firely Server's facade implementers. The Bulk Data Export page now has a section on [BDE for facades](#).
2. Upgraded the .NET SDK to 3.7.0. Please review its [release notes](#) for changes.

#### Logging improvements

1. Error messages including information about authorization validation and authentication requests are now enriched with user information if `ShowAuthorizationPII` is enabled [in the configuration](#).
2. Authorization/Authentication logging messages are now enriched with more information when logging level for the namespace `Vonk.Smart` is set to `Debug`.
3. In case [SSL is activated](#), but the `.pfx` file configured in `CertificateFile` could not be found, Firely Server will now log this error more explicitly.

### Fix

1. Fixed a bug where newly created SQL connections were not closed properly with the raw SQL configuration.
2. Fixed a bug that prevented searching on the ContactPoint datatype with a query of type `system|value`. Although this combination is disallowed by the FHIR specification, Firely Server still allows it. We do not provide a migration for this issue. Please [Contact us](#) if this is an issue for you.
3. Fixed a bug that returned invalid self links without escaped whitespaces in bundles.
4. Improved support for use of Firely Server with Azure SQL.

### Other

1. Firely Server will no longer support CosmosDb starting with version 4.7.0.
2. The Docker image name has changed from `simplifier/vonk` to `firely/server`. The old image name will be maintained for a few months to allow for a smooth transition. When updating to version 4.7.0, you should start to use the new image name. Versions 4.6.2 and older will stay available (only) on 'simplifier/vonk'.

## 3.1.12 Release 4.6.2, Dec 23rd, 2021

### Fix

1. `IConformanceCacheR3` and `IConformanceCacheR4` are registered again in the `ServiceProvider` for plugins that still make use of them. Note that these interfaces are obsolete by now, so make sure you don't use them for any new plugins.

## 3.1.13 Release 4.6.1, Dec 15th, 2021

### Fix

1. Improved handling of `TypeLoadException` and `ReflectionTypeLoadException` when scanning external assemblies for `SerializationSupportAttribute` attributes.

## 3.1.14 Release 4.6.0, Nov 18th, 2021

### Database

1. SQL Server (all changes below applicable only when plugin `Vonk.Repository.Sql.Raw` is enabled)
  1. A new computed column `IsDeleted` on table `[vonk].[entry]` is leveraged for more performant SQL queries

---

**Note:** The performance of the old `Vonk.Repository.Sql` may be adversely impacted by this change. We encourage you to use the new `Vonk.Repository.Sql.Raw` implementation.

---

2. Improved performance of SQL queries by converting 5 columns from `[vonk].[entry]` to `varchar` upon retrieval: `InformationModel`, `Type`, `ResourceId`, `Version`, `Reference`

---

**Note:** These columns should - by definition of the FHIR datatypes - not contain characters outside the varchar range, but please pay attention to this change if your id's or custom resource type has those characters nonetheless. We may alter the datatype of the columns in a future release.

---

3. Improved performance of some SQL queries by avoiding unnecessary SQL query parameter type conversion
  4. Improved performance of some SQL queries by avoiding excessive retrieval of the (large) ResourceJson column
2. MongoDB
1. Improved performance of searches within a compartment
  2. Added an index `ix_sysinfo` to quickly retrieve the `VonkVersion` document.

## Features

1. Added support for SMART on FHIR v2

---

**Note:** Since most users currently use SMART on FHIR v1, the plugin for v2 is by default *disabled* in the PipelineOptions. You can switch v1 out and v2 in when you want to test the use of v2.

---

## Logging improvements

1. The password and the username are stripped out from a connection string when it gets logged (SQL Server / Sqlite, Verbose log level)
2. SQL param values are not logged by default. This can be enabled by using a new config setting. See [SQL query parameter logging](#) (SQL Server / Sqlite, Verbose log level)
3. Username and UserId are included in log and audit entries (when using SoF or another authentication plugin)
4. SQL query duration now gets logged (changed for `Vonk.Repository.Sql.Raw.KSearchConfiguration` plugin; was always available for other repository plugins, Verbose log level)
5. Fixed category names for some log entries to include the fully qualified type of their source. For example, category `MetadataConfiguration` was changed to `Vonk.Core.Metadata.MetadataConfiguration`, and category `BulkDataExportConfiguration` was changed to `Vonk.Plugin.BulkDataExport.BulkDataExportConfiguration`, etc.

## Fix

1. Fixed a bug when validation was not performed on PATCH requests even when the validation level was set to Full
2. Fixed a bug when escaping of the pipe ('|') character was not working as expected for token search parameters
3. Improved error handling when FS tries to load a non-.NET DLL from the plugins directory
4. Fixed a bug (introduced in 4.5.1) when a compartment matches more than 1 Patient
5. Fix: `$validate` checks whether a system parameter is provided
6. Fix: `Vonk.Repository.Sql.Raw`: searching on quantities with values having a high precision failed

### Other

1. Firely SDK upgraded from v3.0.0 to v3.6.0. See the SDK release notes [here](#)

---

**Note:** This will make Firely Server import a new version of specification.zip into the Administration endpoint for each FHIR version. If you share the Administration database among instances, allow 1 instance to finish this process before starting the other instances.

---

### 3.1.15 Release 4.5.1

#### Attention:

The upgrade procedure for Firely Server running on MongoDB will execute an upgrade script that adds a new field to store precalculated compartment links. If your collection contains a lot of resources, this may take a very long time. Therefore, the MongoDB upgrade script has to be executed manually. The script can be found in *mongodbFS\_SchemaUpgrade\_Data\_v17\_v18.js*

Here are some guidelines:

- We tested it on a MongoDB collection with about 400k documents in total. The upgrade script took around 3.5 minutes to complete on a fairly powerful laptop.
- As always, make sure you have a backup of your database that has been tried and tested before you begin the upgrade.
- Please make sure that Firely Server is shutdown before you execute the script.
- If you encounter problems running the script, or need any assistance, please [Contact us](#).

### Database

1. MongoDB
  1. The migration script 'FS\_SchemaUpgrade\_Data\_v17\_v18.js' has been fixed. All data present in the database before the migration is now again accessible after the migration.
2. SQL Server
  1. Improved the query performance when using `_include` by using "WITH FORCESEEK".
  2. Improved performance by avoiding scanning indexes when searching on the UriHash column

### Fix

1. Firely Server will now by default include a user-agent header when retrieving the SMART Discovery document



### 3.1.16 Release 4.5.0

#### Database

**Attention:** The release version of the MongoDB migration contains an error causing compartment searches to return no search results for all migrated resources. Only newly added resources after the migration will be returned successfully. In *Release 4.5.1* we have fixed this issue, so please use that version instead.

##### 1. MongoDB

1. To improve the performance of compartment searches, Firely Server now precalculates the compartment links to which a resource belongs on insert in the database. An external migration script 'FS\_SchemaUpgrade\_Data\_v17\_v18.js' is provided in the distribution. It needs to be applied manually using MongoDB Shell.

#### Security

1. A VonkConfigurationException, which was thrown if a SQL database migration could not be performed, included the SQL connection string in plain text in the log. Please check you log files if they include any sensitive information such as the database password, which might have been part of the connection string.

#### Fix

1. It is now possible to configure pre- and post-handlers for a custom operations using VonkInteraction.all\_custom regardless of the interaction level of the operation handler and the interaction level on which the operation is configured in the appsettings.
2. \$lastN could not handle chained arguments on the subject/patient reference
3. \$lastN reported an invalid error message if the reference to a subject/patient was provided as an urn:uuid reference
4. \$lastN search result bundles were missing self-links when no results were found
5. Disabling Vonk.Fhir.R4 in the pipeline resulted in an internal exception thrown by the ConformanceCache

#### Feature

1. \$lastN can be combined with \_elements and \_include parameters
2. \$lastN can group the results by the component-code or combo-code search parameter

#### Documentation

1. Added an explanation to the documentation why the use of \_total=none influences the performance of a search query.

### Plugins

1. The FHIR Mapper is no longer distributed together with Firely Server. Please contact [fhir@healex.systems](mailto:fhir@healex.systems) for any questions regarding the FHIR Mapper.
2. The packages Vonk.Fhir.R(3|4) depended on an unpublished NuGET package Vonk.Administration.Api.
3. All classes in the namespace 'Vonk.Facade.Relational' are now published on [GitHub](#).

### 3.1.17 Release 4.5.0-beta

#### Fix

1. Security: Added a warning to the documentation that using compartments other than 'Patient' to restrict access based on patient-level SMART on FHIR scopes may result in undesired behavior. See [Compartments](#) for more information.
2. The RequestCountService caused an exception on startup if the RequestInfoFile could not be accessed, e.g. due to limited filesystem permissions. The RequestCountService has been removed completely. Any remaining .vonk-request-info.json files can be deleted manually.
3. The logsettings for SQL server included an outdated configuration.
4. The logsettings for MongoDB included an outdated configuration.

#### Feature

1. Improved error messages if an internal exception occurred due to failing filesystem access.
2. The [\\$lastN](#) operation is now available when using SQL Server as the backend for Firely Server. See [\\$lastN Observations](#) for more information.

### Plugin and Facade

1. Added async support for the ISnapshotGenerator interface and its implementations.

### 3.1.18 Release 4.4.0

#### Database

1. MongoDB
  1. To improve the performance of deletes, the definition of the index `ix_container_id` is redefined. Firely Server 4.4.0 will automatically change the definition.
2. SQL Server
  1. Improved query behind `_include` to leverage an index. No changes to the database schema involved. This only affects the new implementation (available since 4.3.0).

## Fix

1. Improved automatic upgrading of terminology settings from pre-4.1.0 instances.
2. Added `CapabilityStatement.status` for R4
3. The default `SmartAuthorizationOptions` in `appsettings.default.json` only have the Filter for 'Patient' enabled. The rest is now commented out as those are generally not used.

## Plugin and Facade

1. The interfaces `PrioritizedResourceResolver(R3|R4|R5)` and their implementations are no longer available. It is advised to construct your own `StructureDefinitionSummaryProvider` incl. a `MultiResolver` combining your own resource resolver and the `IConformanceCache` provided by Firely Server.
2. The interface `IConformanceCacheInvalidation` has been moved from `Vonk.Core.Import` to `Vonk.Core.Conformance`
3. The classes `SpecificationZipResolver(R3|R4|R5)` are no longer available. Please use the `IPrioritizedResourceResolvers` instead.
4. Starting from this version, a Facade should not have an order greater than or equal to 211. The reason for this is that upon configuring the administration database, Firely Server checks whether an `ISearchRepository` is registered. The earliest of these configurations is at order 211.

## 3.1.19 Release 4.3.0

### Database

1. SQL Server
  1. To improve the performance of searching we have rewritten a large part of our SQL Server implementation. To be able to use the new implementation go to section `PipelineOptions` in `appsettings.default.json` (or `appsettings.instance.json` if you have overridden the default pipeline options) and add `"Vonk.Repository.Sql.Raw.KSearchConfiguration"`. See [Using SQL server](#) for more details.
  2. We have identified two indexes that needed a fix to increase query performance for certain searches. The upgrade procedure will try to fix these indexes automatically. If your database is large, this may take too long and the upgrade process will time out. If that happens you need to run the upgrade script manually. The script can be found in `sqlserver/FS_SchemaUpgrade_Data_v19_v20.sql`. If you use SQL Server as your Administration database, Firely Server will try to update it automatically as well. If you prefer a manual update, you can run the following script: `sqlserver/FS_SchemaUpgrade_Admin_v18_v19.sql`.

### Feature

1. Firely Server now allows you to execute a `ValueSet` expansion of large `ValueSets` (> 500 included concepts). Previously, Firely Server would log an error outlining that the expansion was not possible. The `appsettings` now contain a setting in the Terminology section allowing to select the `MaxExpansionSize`. See [Options](#) for more details.

### Fix

1. Fixed a `NullPointerException` which occurred when indexing UCUM quantities that contained more than one annotation (e.g. “`{reads}/{base}`”).
2. Fixed a bug where it was possible to accidentally delete a resource with a different information model then the request. Firely Server will now check the information model of the request against the information model of the resource for conditional delete and delete requests.
3. `$subsumes` returned HTTP 501 - Not implemented for a POST request (instance-level) even if the operation was enabled in the appsettings.
4. The `_type` filter on `$everything` and Bulk data export didn't allow for resources that are not within the Patient compartment. The operations would return an empty result set.
5. Added a clarification to the documentation that `$everything` and Bulk data export do not export Device resources by default. Even though the resource contains a reference to Patient, the corresponding compartment definition for Patient does not include Device as a linked resource. It is possible to export Device resources by adding the resource type to “AdditionalResources” settings of the operations.

### 3.1.20 Release 4.2.1 hotfix

#### Database

---

**Note:** We found an issue in version 4.2.0, which affects the query performance for Firely Server running on a SQL Server database. If your are running FS v4.2.0 on SQL Server you should upgrade to v4.2.1 or if that is not possible, [Contact us](#).

---

#### Attention:

The upgrade procedure will execute a SQL script try to validate the foreign key constraints. If your database is large, this may take too long and the upgrade process will time out. If that happens you need to run the upgrade script manually, The script can be found in `data/20210720085032_EnableCheckConstraintForForeignKey.sql`.

Here are some guidelines:

- We tested it on a database with about 15k Patient records, and 14 million resources in total. The upgrade script took about 20 seconds to complete on a fairly powerful laptop.
- As always, make sure you have a backup of your database that has been tried and tested before you begin the upgrade.
- If you expect the upgrade to time out, you can choose to run the SQL script manually beforehand. Please make sure that Firely Server is shutdown before you execute the script.

## Fix

1. Fixed a bug where some of the Foreign Keys in SQL Server had become untrusted. This bug has an impact on the query performance since the the SQL Server query optimizer will not consider FKs when they are not trusted. This has been fixed, all Foreign Keys have been validated and are trusted again.

### 3.1.21 Release 4.2.0

#### Database

**Attention:** For SQL Server users: this version of Firely Server running on SQL Server has a bug where some of the Foreign Keys became untrusted. This has an impact on the query performance. Please upgrade to version 4.2.1 or if that is not possible, [Contact us](#). Please note that users running Firely Server running either MongoDB, CosmoDb, or SQLite are not affected by this issue.

**Attention:** For SQL Server we changed the datatype of the primary keys. The related upgrade script (`data/20210519072216_ChangePrimaryKeyTypeFromIntToBigint.sql`) can take a lot of time if you have many resources loaded in your database. Therefore some guidelines:

- We tested it on a database with about 15k Patient records, and 14 mln resources in total. Migrating that took about 50 minutes on a fairly powerful laptop.
- Absolutely make sure you create a backup of your database first.
- If you haven't done so already, first upgrade to version 4.1.x.
- If you already expect the migration might time out, you can run it manually upfront. Shut down Firely Server, so no other users are using the database, and then run the script from SQL Server Management Studio (or a similar tool).
- Running the second script (`20210520102224_ChangePrimaryKeyTypeFromIntToBigintBDE.sql`) is optional - that should also succeed when applied by the auto-migration.

## Feature

1. Terminology operation \$lookup is now also connected to remote terminology services, if enabled. See [Terminology services](#).
2. We provided a script to 'purge' data from a SQL Server database. See `data/20210512_Purge.sql`. You can filter on the resource type only. Use with care and after a backup. If you need more elaborate support for hard deletes, please [Contact us](#).

### Fix

1. Firely Server could run out of primary keys on the index tables in SQL Server. Fixed by upgrading to bigint, see warning above.
2. Nicer handling of SQL Server migration scripts that time out on startup. It will now kindly ask you to run the related script manually if needed (usually depends on the size of your database).
3. The Patient-everything (\$everything) operation was not mentioned in the CapabilityStatement.
4. License expired one day too early.
5. Dependencies have been upgraded to the latest versions compatible with .NET Core 3.1.
6. PATCH did not allow adding to a repeating element.
7. If your license does not allow usage of SMART on FHIR, authorization was disabled, emitting a warning in the log. Possibly causing unauthorized access without the administrator noticing it. This specific case will now block the startup of Firely Server.

### 3.1.22 Release 4.1.3 hotfix

#### Fix

1. Fixed a bug where a number of concurrent \$transform requests on a freshly started Firely Server could lead to Internal Server Error responses.
2. Upgraded the Mapping plugin.

### 3.1.23 Release 4.1.2 hotfix

#### Fix

1. Fixed a bug when trying to delete multiple resources at once (bulk delete, see [Configuration](#) for configuration options). The operation would take a while and eventually return a 204 No Content without actually deleting any resources. This is fixed, the bulk delete operation now deletes the resources.

### 3.1.24 Release 4.1.1 hotfix

#### Feature

1. SMART configuration: Some identity providers use multiple endpoints with different base addresses for its authorization operations. Added an extra configuration option `AdditionalEndpointBaseAddresses` to define additional base endpoints addresses next to the main authority endpoint to accommodate this. See [Configuration](#) for further details.

## Fix

1. Fixed an error in SQL script `data/20210226200007_UpdateIndexesTokenAndDatetime_Up.sql` that is used when manually updating the database to v4.1.0. We also made the script more robust by checking if the current version the database is suitable for the manual upgrade.

### 3.1.25 Release 4.1.0

**Attention:** We have found an issue with SMART on FHIR and searching with `_(rev)include`. And fixed it right away, see Fix nr 1 below. Your Firely Server might be affected if:

- you enabled SMART on FHIR
- and used `patient/read.*` scopes together with a patient compartment

What happens? Patient A searches Firely Server with a patient launch scope that limits him to his own compartment. If any of the resources in his compartment links to *another* patient (let's say for Observation X, the performer is Patient B), Patient A could get to Patient B with `<base>/Observation?_include=Observation.performer`. If you host Group or List resources on your server, a `_revinclude` on those might give access to other Patient resources within the same Group or List.

If you think you might be affected you can:

- upgrade to version 4.1.0
- or if that is not possible, [Contact us](#).

## Database

### 1. SQL Server

1. A new index table was added. The upgrade procedure will try to fill this table based on existing data. If your database is large, this may take too long and time out. Then you need to run the upgrade script found in `data/20210303100326_AddCompartmentComponentTable.sql` manually.
2. A new SQL Server index was added to improve query times when searching with date parameters. The upgrade procedure will try to build this index. If your database is large, this may take too long and time out. Then you need to run the upgrade script found in `data/20210226200007_UpdateIndexesTokenAndDatetime_Up.sql` manually.
3. In both cases you may also run the script manually beforehand.
4. As always: make sure you have a backup of your database that is tested for restore as well.

## DevOps

**Attention:** Because of a change in the devops pipeline there is no longer a `Firely.Server.exe` (formerly `Vonk.Server.exe`) in the distribution zip file. You can run the server as always with `dotnet ./Firely.Server.dll`

## Features

1. Inferno, The ONC test tool: Firely Server now passes all the tests in this suite! With version 4.1.0 we specifically added features to pass the 'Multi-patient API' tests. Do you want a demo of this? [Contact us!](#)

2. Terminology support has been revamped. Previously you needed to choose between using the terminology services internal to Firely Server *or* external terminology services like from OntoServer or Loinc. With this version you can use both, and based on the codesystem or valueset involved the preferred terminology service is selected and queried.
  1. This works for terminology operations like `$validate-code` and `$lookup`
  2. It also works for validation, both explicitly with `$validate` and implicitly, when validating resources sent to Firely Server.
  3. The CodeSystem, ValueSet and ConceptMap resources involved are conformance resources and therefore always retrieved from the Administration database.
  4. Responses may differ on details from previous versions of Firely Server, but still conform to the specification.
  5. See [Terminology services](#) for further details.
3. `$everything`: We now support the *Patient \$everything* operation for single Patients. (For multiple patients, there is the Bulk Data Export feature.)
4. Performance of `$everything`, Bulk Data Export and authorization on compartments improved. We added a special index to the database that keeps track which resource belongs to which compartment. First in SQL Server, MongoDB has less need for it.
5. SMART on FHIR: Support for token revocation. Reference tokens can be revoked, and Firely Server can check for the revocation.

### Fixes

1. SMART on FHIR: We have found ourselves that the authorization restrictions were bypassed when using `_include` or `_revinclude` in a FHIR Search. We solved this security issue immediately.
2. Firely Server transparently translates absolute urls to relative urls (for internal storage) and back. There was a performance gain to be made in this part, which we did. This is mostly notable on large transaction or batch bundles.
3. Batch bundles are not allowed to have links between the resources in the entries. Firely Server will now reject batch bundles that have these links. If you need links, use a transaction bundle instead.

### Plugin and Facade

1. We upgraded the Firely .NET SDK to version **3.0.0**. This SDK version is almost fully compatible with 2.9, but it brings significant simplifications to its use because the Parameters and OperationOutcome resource POCOs are no longer FHIR-version specific.

---

**Note:** Every new version of the SDK brings new versions of the `specification.zip` files. So upon upgrade these new files will be read into the Administration database. See [Managing Conformance Resources](#) for more background.

---



### 3.1.26 Release 4.0.0

This major version introduces a new name: **Firely Server** instead of Vonk. Other than that, this release contains some significant code changes, which could impact you if you run Firely Server with your own plugins.

#### Features

1. Name change Vonk -> Firely Server:
  1. The main entry point dll (formerly: `Vonk.Server.dll`) and executable (formerly: `Vonk.Server.exe`) names have been changed to `Firely.Server.dll` and `Firely.Server.exe` respectively.
  2. The name was changed in the `CapabilityStatement.name`.
  3. The name of the download zip (from Simplifier) has changed from `vonk_distribution.zip` to `firely-server-latest.zip`. Likewise the versioned zip files have changed as well.
2. We have implemented FHIR Bulk Data Access (`$export`) to allow for fast, asynchronous ndjson data exports. The [Bulk Data Export documentation](#) can help you to get started.
3. Firely Server now uses Firely .NET SDK 2.0.2 (formerly: FHIR .NET API)

**Attention:** If you are running Firely Server with your own self-made plugins, you will likely encounter package versioning problems and need to upgrade your NuGet Firely Server package references (package names starting with `Vonk.`) to version 4.0.0. You also need to upgrade any Firely .NET SDK package references (package names starting with `Hl7.Fhir.`) to version 2.0.2. The [Firely .NET SDK release notes](#) and [Breaking changes in Firely SDK 2.0](#) can give you an idea of the changes you may encounter in the SDK.

4. SMART on FHIR can now recognize prefixes to the claims, see its [Configuration](#).
5. The smart-configuration endpoint (`<url>/.well-known/smart-configuration`) relays the signature algorithms configured in the authorization server.

#### Fixes

1. Application Insights has now been disabled by default. If you need Application Insights, you can enable it in your log settings file by including the entire section mentioned in [Application Insights log settings](#).
2. When validating a resource, a non-existing code would lead to an `OperationOutcome.issue` with the code `code-invalid`. That issue code has been changed to `not-supported`.
3. On a batch or transaction bundle errors were not reported clearly if the entry in error had no `fullUrl` element. We fixed this by referring to the index of the entry in the entry array, and the resource type of the resource in the entry (if any).
4. The `import[R4]` folder allows for importing custom `StructureDefinition` resources. If any of them had no id, the error on that caused an exception. Fixed that.
5. If a Facade returned a resource without an id from the Create method, an error was caused by a log statement. Fixed that.
6. Indexing `Subscription.channel[0].endpoint[0]` failed for R4. Fixed that. This means you can't search for existing Subscriptions by `Subscription.url` on the `/administration` endpoint for FHIR R4.
7. Postman was updated w.r.t. acquiring tokens. We adjusted the [documentation on that](#) accordingly.

8. If a patient claim was included in a SMART on FHIR access token, the request would be scoped to the Patient compartment regardless of the scope claims. We fixed this by allowing “user” scopes to access FHIR resources outside of the Patient compartment regardless of the patient claim. See [Launch context arrives with your access\\_token](#) for more background information.

### Plugin and Facade

1. The mapping plugin is upgraded to the Mapping Engine 0.6.0.
2. As announced in [Release 3.0.0](#) we removed support for creating a Facade as a standalone ASP.Net Core project. You can now only build a Facade as a plugin to Firely Server. See [Firely Server Facade](#) on how to do that.
3. The order of some plugins has changed. This way it is possible to add a plugin between PreValidation and UrlMapping:
  - [UrlMapping](#): from 1230 to 1235
  - [Prevalidation](#): from 4320 to 1228
4. A Facade based on `Vonk.Facade.Relational` no longer defaults to STU3

**Attention:** If you developed a facade plugin based on `Vonk.Facade.Relational`, you need to override `RelationalQueryFactory.EntryInformationModel(string informationModel)` in your implementation to allow the FHIR version you wish to target (see [Deciding on a FHIR version](#))

5. We took the opportunity of a major version upgrade to clean up a list of items that had been declared `Obsolete` already. Others have become obsolete now. This is the full list:

# Obsolete, now deleted:

```
# Vonk.Core.Common.DeletedResource # Vonk.Core.Common.IResource.Currency, Change and
Clone(), also in VonkResource. # Vonk.Core.Common.IResourceExtensions.ToIResource(this
ISourceNode original, ResourceChange change, ResourceCurrency cur-
rency = ResourceCurrency.Current) (the overload defaulting to STU3) #
Vonk.Core.Context.Guards.SupportedInteractionOptions.SupportsCustomOperationOnLevel()
# Vonk.Core.Context.Internal.BatchOptions # Vonk.Core.Operations.Validation.ValidationOptions #
Vonk.Core.Pluggability.InteractionHandlerAttribute.Tag # Vonk.Core.Pluggability.ModelOptions #
Vonk.Core.Repository.SearchOptions.LatestOne # Vonk.Core.Support.LogHelpers.TryGetTelemetryClient,
both overloads. # Vonk.Core.Support.SpecificationZipLocator.ctor(IHostingEnvironment...) #
Vonk.Fhir.R3.IResourceVisitor + extensions # Vonk.Fhir.R3.Configuration.ModelContributorsFacadeConfiguration
# Vonk.Fhir.R3.FhirExtensions.AsIResource() # Vonk.Fhir.R3.FhirPropertyIndex + FhirProp-
ertyInfo + FhirPropertyIndexBuilder # Vonk.Fhir.R3.IConformanceBuilder + BaseC-
onformanceBuilder + HarvestingConformanceBuilder + extensions + IConformance-
Contributor # Vonk.Fhir.R3.CompartmentDefinitionLoader + (I)SearchParameterLoader
# Vonk.Fhir.R3.MetadataImportOptions + MetadataImportSet + ImportSource #
Vonk.Fhir.R3.PocoResource + PocoResourceVisitor # Vonk.Core.InformationModelAttribute
(actually made internal)
```

# Obsolete since this version:

```
# Vonk.Core.Configuration.CoreConfiguration: allows for integrating Vonk components
in your own ASP.NET Web server, discouraged per 3.0 (see these releasenotes). #
Vonk.Fhir.R3.FhirR3FacadeConfiguration: see above.
```

## 3.2 Old Firely Server release notes (v3.x)

### 3.2.1 Release 3.9.3 hotfix

**Attention:** We changed the behavior of resthook notifications on Subscriptions. See Fix nr 1 below.

#### Database

1. SQL Server: The migration that adds the indexes described in [Release 3.9.2 hotfix](#) might run longer than the timeout period of 30 seconds. Therefore we added scripts to apply and revert this migration manually. If you encounter the timeout during upgrade: shut down vonk, run the script using SQL Server Management Studio or any similar tool, then start Vonk 3.9.3 again. In both scripts you only need to provide the database name for the database that you want to upgrade. If you run your administration database on SQL Server you can but probably do not need to run this script on it. The administration database is typically small enough to complete the script within 30 seconds.

1. apply: <vonk-dir>/data/2021211113200\_AddIndex\_ForCountAndUpdateCurrent\_Up.sql
2. revert: <vonk-dir>/data/2021211113200\_AddIndex\_ForCountAndUpdateCurrent\_Down.sql

#### Fix

1. *Subscriptions*: A resthook notification was sent as a FHIR create operation, using POST. This was not compliant with the specification that states it must be an update, using PUT. We changed the default behavior to align with the specification. In order to avoid breaking changes in an existing deployments, you may set the new setting `SubscriptionEvaluatorOptions:SendRestHookAsCreate` to `true` - that way Vonk will retain the (incorrect) behavior from the previous versions.

### 3.2.2 Release 3.9.2 hotfix

#### Fix

All fixes are relevant to SQL Server only.

1. The 3.9.0 fix that “Improved the handling of concurrent updates on the same resource.” decreased the performance of concurrent transaction handling. We implemented another solution that does not affect performance.
2. Improved read performance by adding an index.

### 3.2.3 Release 3.9.1 hotfix

#### Fix

1. Fixed a bug introduced with 3.9.0 where Vonk would throw the following exception on start-up: `System.InvalidOperationException: Unable to resolve service for type 'Vonk.Core.Conformance.IDefinitionProvider' while attempting to activate 'Vonk.Fhir.R3.SnapshotGeneration.SnapshotGeneratorR3`
2. Fixed a breaking change to public search API with the implementation of `_total` parameter. We had introduced a new parameter to the Next method in ResultPage, effectively breaking backwards compatibility. This has been fixed.

### 3.2.4 Release 3.9.0

#### Features

1. We have made Subscriptions more robust. See [Subscriptions](#) for details. In summary, if an evaluation of a Subscription fails, Vonk will retry the evaluation periodically for a number amount of tries. You can control the retry period and the maximum number of retries in the subscription settings:
  - `RetryPeriod` is expressed in milliseconds. Default `300000` (30 sec).
  - `MaximumRetries` is the maximum amount of times Vonk will retry to send the resources. Default 3 retries.
2. We have implemented the `_total` parameter for options `none` and `accurate`. Omitting the `_total` parameter is equivalent to `_total=accurate`. If the total number of resources is not relevant, using `_total=none` in the request results in better performance when searching.
3. It is no longer necessary for the `:type` parameter to always be provided to distinguish between multiple reference targets. The parameter does not need to be provided anymore when the search only applies to a single target. For example: `GET <base>/AllergyIntolerance?patient=xyz` The `patient:Patient` type parameter does not have to be supplied. The ‘patient’ search parameter on `AllergyIntolerance` has two possible targets. It may reference either a `Patient` or a `Group` resource. However, the fhirpath statement that goes with it, selects ‘`AllergyIntolerance.patient`’, and that reference element may only target a `Patient` resource.

#### Fixes

1. Indexing values for a string search parameter threw an exception when there was no value but only an extension. This has been corrected.
2. We made the `Provenance.target` available as a `revInclude` Parameter in the `CapabilityStatement`. Previously, Vonk did not account for the case that a reference is allowed to ANY resource type, which incorrectly resulted in `Provenance.target` to not be shown in the `CapabilityStatement`.

All the following fixes are only relevant for SQL Server:

1. Improved the handling of concurrent updates on the same resource.
2. Upgraded the version of the `SqlClient` library to fix issues when running on Linux.
3. Fixed missing language libraries for SQL Server when running on Docker.

### 3.2.5 Release 3.8.0

#### Database

- We added an important note to the [3.6.0 release notes](#) for MongoDB users.
- Because of the changes in searching for Quantities (feature 2 below), you will need to do a [reindex](#) in order to make use of this. You may limit the reindex to only the search parameters of type ‘quantity’ that you actually use (e.g. `Observation.value-quantity`).

## Features

1. We upgraded the FHIR .NET API to 1.9, see the [1.9 releasenotes](#). This will trigger an automatic *import of the Conformance Resources* at startup.
2. We upgraded the [Fhir.Metrics library](#) to 1.2. This allows for a more uniform search on Quantities (mainly under the hood)
3. We upgraded the FHIR Mapping plugin to support the FHIR Mapper version 0.5.
4. The *built-in terminology services* now support the `includeDesignations` parameter.
5. The *IVonkContext* now lets you access the original `HttpContext`.
6. The `CapabilityStatement` now lists the profiles that are known to Vonk (in its Administration database) under `CapabilityStatement.rest.resource.supportedProfile` ( $\geq$  R4 only) and the base profile for a resource under `CapabilityStatement.rest.resource.profile`.
7. We extended the security extension on the `CapabilityStatement` to include the endpoints for `register`, `manage`, `introspect` and `revoke`.
8. `IAdministrationSearchRepository` and `IAdministrationChangeRepository` interfaces are no publicly available. Use with care.

## Fixes

1. If the server configured as authorization endpoint in the Smart setting is not reachable, Vonk will log a proper error about that.
2. An error message for when a query argument has no value is improved.
3. When *SMART-on-FHIR* is enabled, and the received token contains a launch context, the *\_history* operation is no longer available. Because Vonk does not retain the search parameter index for historical resources, it cannot guarantee that these resources fall within the launch context (at least not in a performant way). To avoid information leakage we decided to disable this case altogether.
4. A Create interaction without an id in the resource, with *SMART-on-FHIR* enabled, resulted in an exception.
5. You can now escape the question mark '?' in a query argument by prepending it with a backslash '\'
6. A Quantity search using 'lt' on MongoDB resulted in too many results.

## 3.2.6 Release 3.7.0

### Database

**Attention:** To accommodate for feature #2 below there is an automatic migration carried out for SQL Server and SQLite. This migration might take some time, so please test it first. For MongoDB, you will have to *Rebuild the whole search index*. If this is not feasible for your database, please *Contact us* for assistance.

## Features

1. Patch: We implemented FHIR Patch. You can now update a resource having only partial data for it. See *Create, read, update, patch, delete*.
2. Search on accents and combined characters: we improved searching with and on accents and combined characters. Note the database change above.

3. API 1.7: We upgraded Vonk to use the FHIR .NET API 1.7, having its own [releasenotes](#).
4. Security: The Docker image is now based on the Alpine image for .NET Core. This has far less security issues than the Ubuntu image that we used before. The base image is aspnet:3.1-alpine:3.11 (newest version 3.12 has an open bug related to SQLite).
5. Security: We revisited the list of security vulnerabilities, see [Security notifications for Firely Server](#).
6. Administration: ConceptMaps are now *imported* at startup.

### Fixes

1. Searching on `_lastUpdated` could be inaccurate when time zone differences are in play. We fixed that.
2. Search arguments for a quantity search weren't allowed to be greater than 999.

## 3.2.7 Release 3.6.1

### Features

1. ConceptMap resources can be stored at the Administration endpoint, both through *import* and through the *RESTful API*.

**Attention:** Previous versions of Vonk did not include the ConceptMap resources in the import so they will currently not be in your Administration database. If you run your Administration database on SQL Server or MongoDB and want to use the ConceptMap resources from the spec, be sure to rerun the import of the specification resources. You can force Vonk to do so by deleting the `.vonk-import-history.json` file from the ImportedDirectory (see the settings under `AdministrationImportOptions`). If you use SQLite, you can simply use the pre-built `./data/admin.db` from the binaries.

### Plugins

1. The FHIR Mapper plugin is upgraded to version 0.3.6.
2. The FHIR Mapper plugin now fully works on the *Administration* endpoint.

## 3.2.8 Release 3.6.0

### Database

**Attention:** For MongoDB users: We implemented feature #1 below using the Aggregation Pipeline. This makes an existing issue in MongoDB - *SERVER-7568* <<https://jira.mongodb.org/browse/SERVER-7568>> - a more urgent problem. MongoDB has solved this problem in version 4.4. Therefore we advise you to upgrade to MongoDB 4.4.

### Feature

1. Sort: The *sorting* that was implemented for the SQL/SQLite repositories in the previous version is now also implemented for MongoDB.

2. Terminology: The *local terminology service*, built in to the Vonk Administration API, is upgraded to support R4 and R5 (and still R3 of course).
3. Vonk can now index and search on search parameters that reference a nested resource, like Bundle.message.

**Attention:** Note that any nested resources have to be indexed by Vonk. For new data that is done automatically. But if you want to use this on existing data, you have to *reindex for the search parameters* you want to use it on. Those will most notably be Bundle.message and Bundle.composition.

4. If you accidentally provide a body in a GET or DELETE request, Vonk will now ignore that body instead of returning an error.

## Fix

1. CapabilityStatement (rev)includes now use ‘:’ as a separator instead of ‘.’.

## Plugins

1. The *BinaryWrapper plugin* is upgraded to 0.3.1, where the included BinaryEncodeService is made more reusable for other plugins (most notably the FHIR mapper).

## 3.2.9 Release 3.5.0

### Feature

1. Search reference by identifier: FHIR R4 allows you to *search a reference by its identifier*. We added support for this in Vonk. Note that any identifiers in reference elements have to be indexed by Vonk. For new data that is done automatically. But if you want to use this on existing data, you have to *reindex for the search parameters* you want to use it on. E.g. Observation.patient.
2. AuditEvent logging: In release 3.3.0 we already added support for logging audit information to a file. With this release we add to that logging that same information in AuditEvent resources. These resources are written to the Vonk Data database (not the Administration database). Users are not allowed to update or delete these resources. See *Auditing* for more background.
3. Audit logging: We added [Request] or [Response] to the log lines so you can distinguish them better.
4. Sort: We started implementing *sorting*. This release provides sorting for search parameters of the types string, number, uri, reference, datetime and token, on the repositories SQL, SQLite and Memory. On the roadmap is extending this support to MongoDB and to quantity search parameters.
5. *Terminology Integration*: You can configure Vonk to route the terminology operations to external terminology servers. You can even configure a preferred server for certain code systems like LOINC or Snomed-CT. On the roadmap is to also allow you to use these servers for validation of codes and for token searches.
6. We implemented *\$meta-delete*, see *Meta plugins*.
7. Loading plugins can lead to unexpected errors. We made the process and the log friendlier, so you can spot configuration errors more easily:
  - The log outputs the version of each of the plugins
  - If a duplicate .dll file is found, Vonk tells you which two dlls are causing this and then exits.

- If you configured a plugin that you are not licensed to use, this is logged with a friendly hint to acquire a license that does allow you to use it.
8. The log is now by default configured to use asynchronous logging so Vonk is not limited by the speed of the logging sinks (like the Console and the log file). Please update your `logsettings.instance.json` if you created your own log settings in that. See [Log settings](#) for more background.

### Fix

1. You could load invalid XML in the `Resource.text` through a JSON payload. When that resource was then retrieved in XML, it would fail with an `InternalServerError`. Vonk will now return an `OperationOutcome` telling you what the problem is. You can then correct it by using JSON.
2. Composite search parameters were not parsed correctly. Now they are. So you don't see warnings like `Composite SearchParameter 'CodeSystem.context-type-quantity' doesn't have components.` anymore.
3. Indexing for the `_profile` search parameter was broken for R4 since Vonk 3.2.1. We fixed it. If you added new resources with Vonk 3.2.1 - 3.4.0, you need to [reindex for the Resource.\\_profile](#) parameter.
4. Audit log: `%temp%` in the path setting was evaluated as `<current directory>\%temp%`. Fixed that to evaluate to the systems temporary directory.
5. The `logsettings.json` configured the Serilog RollingFile sink by default. That is deprecated, so we replaced it with the File sink.
6. [Rebuild the search index for specific searchparameters](#) now returns an error if you forget to actually specify a search parameter.
7. An `InternalServerError` was returned when you validate a resource that is syntactically incorrect. Like a Patient with multiple id's. Vonk now returns an `OperationOutcome` with the actual problem.
8. The configuration for the FHIR Mapper was simplified. You only need to include `Vonk.Plugin.Mapping`. Check `appsettings.default.json` for the new pipeline.
9. Maybe you got accustomed to ignoring a list of warnings at startup of Vonk. We cleaned up the list so that if there is a warning, it is worthwhile investigating the cause of it.
10. The `appsettings` and `logsettings` can contain relative file paths for several settings, like the `License:LicenseFile`. These were evaluated against the current working directory, but that could lead to problems if that was *not* the Vonk directory. We solved that: all relative paths are evaluated against the Vonk directory.
11. The docker image for version 3.4.0 was tagged `3.4.0-`. With 3.5.0 we removed the superfluous hyphen at the end.
12. We updated the documentation on [Using Firely Server on Docker](#) on SQL Server to be clearer about the order of the steps to take.
13. We updated the documentation on [Firely Server Plugin example - Create a new landing page](#) to match .NET Core 3.1.



## Plugins & Facade

1. FHIR Mapper
  - Has been upgraded to version 0.3.4.

### 3.2.10 Release 3.4.0

#### Feature

1. Upgraded to FHIR .NET API 1.6.0, that features a couple of changes for working with CDA logical models. See the [release notes of the API](#).
2. Included the FHIR Mapper in the distribution. It is only enabled however when you include the mapping plugin in your license.

#### Fix

1. When prevalidation is set to the level 'Core', Vonk no longer complains about extensions that are not known if they are not core extensions (i.e. having a url starting with '<http://hl7.org/fhir/StructureDefinition/>').

### 3.2.11 Release 3.3.0

**Attention:** To use the new features for auditing and R5, you need a new license file including the tokens for those plugins. For evaluation and community editions you can retrieve them from Simplifier.net. If you need these updates in your production license, please contact us.

#### Feature

1. Vonk was upgraded to FHIR .NET API 1.5.0. See the [release notes of the API](#).
2. Vonk can now log audit lines in a separate file. This can help you achieve HIPAA/GDPR compliancy. See [Auditing](#) for more info.
3. Failed authorization attempts are now logged from the [SMART on FHIR](#) plugin.
4. Support for `_include:iterate` and `_revinclude:iterate`, see [Search](#).
5. The [BinaryWrapper plugin](#) is now two-way. So you can POST binary content and have it stored as a Binary resource, and GET a Binary resource and have it returned in its original binary format.
6. Experimental support for R5 is now included in the Vonk distribution. For enabling it, see [Support for R5 \(experimental!\)](#).

### Fix

1. Indexing of a quantity in resource could fail with a Status code 500 if it had no `.value` but only extensions.
2. The use of a SearchParameter of type `reference` having no `target` failed. These search parameters are now signalled upon import.
3. Since R4 it is valid to search for a quantity without specifying the unit. Vonk now accepts that.
4. A transaction response bundle could contain an empty `response.etag` element, which is invalid.
5. *Preloading resources* was not working since the upgrade to .NET Core 3.0. That has been fixed. It is still only available for STU3 though.
6. Administration import would state that it moves a file to history when it had imported it. That is no longer true, so we removed this incorrect statement from the log.
7. `$validate-code` could cause a `NullReference` exception in some case.
8. The generated CapabilityStatement for R4 failed constraint `cpb-14`.
9. Content negotiation favoured a mediatype with quality `< 1` over a mediatype without quality. But the default value is 1, so the latter is now favoured.
10. *Validate an instance from the database* did not account for the `informationmodel` (aka FHIR version) of the resource.

### Plugins & Facade

1. *Document Operation*
  - Has been upgraded to Vonk 3.2.0.
  - Was assigned a license token
  - Assigns an id and lastUpdated to the result bundle
2. *Conversion*
  - Has been upgraded to Vonk 3.2.0.
  - Was assigned a license token.
3. `Vonk.Facade.Starter` has been upgraded to Vonk 3.2.1 and as a consequence also to EntityFrameworkCore 3.1.0.

### 3.2.12 Release 3.2.1

### Fix

1. SMART plugin now understands multiple scopes per access token.
2. SMART plugin now understands `Resource.*` claims, in addition to already understanding `Resource.read` and `Resource.write`.

### 3.2.13 Release 3.2.0

**Attention:** Vonk 3.2.0 is upgraded to .NET Core 3.1.0, ASP.NET Core 3.1.0 and EntityFramework Core 3.1.0.

- For running the server: install the ASP.NET Core runtime 3.1.0.
- For developing or upgrading Facades that use Vonk.Facade.Relational: upgrade to EF Core 3.1.0.
- Plugins that target NetStandard 2.0 need not be upgraded.

#### Database

1. There are no changes to the databases. The upgrade of EntityFramework Core does not affect the structure of the SQL Server or SQLite databases, just the access to it.

#### Fix

1. *Supported interactions* were not enforced for custom operations like e.g. \$convert.
2. If a resource failed *Validating incoming resources*, the OperationOutcome also contained issues on not supported arguments.
3. A search with ?summary=count failed.
4. Added support for FhirPath hasValue() method.
5. Resolution of canonical `http://hl7.org/fhir/v/0360|2.7` failed.
6. CapabilityStatement.rest.resource.searchInclude used ‘.’ as separator, fixed to use ‘:’ in <resource>:<search parameter code>
7. Changed default value of License:LicenseFile to vonk-license.json, aligned with the default naming if you download a license from Simplifier.
8. *Reindexing* always interpreted a resource as STU3. Now it correctly honours the actual FHIR version of the resource.

#### Feature

1. *BinaryWrapper plugin* can now be restricted to a list of mediatypes on which to act.
2. Vonk used to sort on \_lastUpdated by default, and add this as extra sort argument if it was not in the request yet. Now you can configure the element to sort on by default in BundleOptions.DefaultSort. Although Vonk FHIR Server does not yet support sorting on other elements, this is useful for Facade implementations that may support that (and possibly not support sort on \_lastUpdated). See also *Search size*.
3. Implemented \$versions operation
4. Extended the documentation on:
  - *The order of plugins*
  - *Constructing bundles*
  - several smaller additions

5. The SMART authorization plugin can now be configured to *not* check the audience. Although not recommended, it can be useful in testing scenarios or a highly trusted environment.

**Attention:** We changed the default value for the setting `SmartAuthorizationOptions.Audience` from `vonk` to empty, or ‘not set’. This is to avoid awkward syntax to override it with ‘not-set’. But if you rely on the value `vonk`, please override this setting in your `appsettings[.instance].json` or environment variables as described in [Changing the settings](#).

### Plugin and Facade API

1. `Vonk.Facade.Relational` now supports the use of the `.Include()` function of EntityFramework Core. To do so, override `RelationalQuery.GetEntitySet(DbContext dbContext)`.
2. `Vonk.Facade.Relational` now supports sorting. Override `RelationalQueryFactory.AddResultShape(SortShape sortShape)` and return a `RelationalSortShape` using the extension method `SortQuery()`.

### 3.2.14 Release 3.1.3 hotfix

#### Fix

1. Fixed behavior on conditional updates in transactions. In odd circumstances Vonk could crash on this.

### 3.2.15 Release 3.1.0

Please also note the changes in [3.0.0](#) (especially the one regarding the SQL server database)

#### Fix

1. Validation on multi-level profiled resources no longer fails with the message “*Cannot walk into unknown StructureDefinition with canonical*”
2. Improved documentation on [upgrading Vonk](#), the [Vonk pipeline](#), [CORS support](#), [plugins](#) and [IIS deployment](#)
3. Using multiple parameters in `_sort` led to an error for all repositories
4. Vonk UI capability statement view now works for self-mapped endpoints like `/R3` or `/R4`
5. A saved resource reference (e.g. `Patient.generalPractitioner`) on a self-mapped endpoint (e.g. `/R4/...`) would have its relative path duplicated (`/R4/R4/...`)

**Attention:** If you have used [self-mapped endpoints](#) (appsettings: `InformationModel.Mapping.Map` in the ‘Path’ mapping mode) and you have saved resources containing references, it is possible that your database now contains some resources with broken references. Please contact us if this is the case

## Feature

1. The new experimental FHIR mapping engine, which is currently exclusively available on our public FHIR server <http://vonk.fire.ly>
2. New licensing system, supporting the *community edition*
3. Simplifier projects are now imported for FHIR R4 as well
4. The following plugins have been bundled with the Vonk release (compare your appsettings with the new appsettings.default.json to activate them)
  1. The \$document operation (see [Documents](#))
  2. The \$convert operation (see [Conversion](#))
  3. The binary wrapper (see [Binary](#))

## Plugin and Facade API

1. Vonk.Facade.Starter has been upgraded to work with Vonk 3.1.0
2. IConformanceContributor and IConformanceBuilder have moved from Vonk.Core.Pluggability to Vonk.Fhir.R3.Metadata. It is also deprecated, as Vonk.Core.Metadata.ICapabilityStatementContributor is now preferred instead. See [Capabilities](#) for more information
3. Implementations of ISearchRepository can now sort on multiple parameters (in BaseQuery.Shapes). Previously this would result in an error.
4. Improved documentation on the [Important classes and interfaces](#)
5. See [Release 3.0.0](#) for some additional issues you may encounter upgrading your plugins

## 3.2.16 Release 3.0.0

### Database

Please also note the changes in [3.0.0-beta1](#)

1. SQL Server: SQL script '20190919000000\_Cluster\_Indexes\_On\_EntryId.sql' (found in the /data folder of the Vonk distribution) must be applied to existing Vonk SQL databases (both to the admin and to the data repositories)

**Attention:** Vonk 3.0.0 (using SQL server) will not start unless this script has been applied to the databases. Please note that running the script can take considerable time, especially for large databases.

## Feature

1. Information model (= FHIR version) settings
  1. Although Vonk now supports multiple information models (STU3 and R4) simultaneously, an unused model can be disabled (see [Configuring the Firely Server Pipeline](#))
  2. You can set the default (or fallback) information model (previously: STU3), which is used when Vonk can not determine the information model from context (see [Information model](#))
  3. You can map a path or a subdomain to a specific information model (see [Information model](#)), mitigating the need to specify it explicitly in a request

2. Vonk now uses [FHIR .NET API 1.4.0](#)
3. Several performance enhancements have been made for SQL server and IIS setups
4. Added R4-style [Conditional Update](#) to both STU3 and R4

### Fix

1. Circular references within resources are now detected, cancelling validation for now. We will re-enable validation for these resources when the FHIR .NET API has been updated
2. An \$expand using incorrect data returned a 500 (instead of the correct 400)
3. Vonk now returns a 406 (Not Acceptable) when the Accept header contains an unsupported format
4. Deletes did not work for R4
5. Search parameters
  1. Search parameters were read twice (at startup and upon the first request)
  2. Search parameter 'CommunicationRequest.occurrence' is not correctly specified in the specification. We provide a correct version.
6. `_history`
  1. `_history` was not usable in a multi information model setup
  2. The resulting `Bundle.entry` in an STU3 `_history` response contained the unallowed response field
  3. Added `Bundle.entry.response` to the R4 `_history` entry
7. Batches
  1. Valid entries in batches also containing invalid entries were not processed
  2. Duplicate `fullUrls` are no longer accepted in a batch request, which previously led to a processing error
  3. An R4 transaction resulted in STU3 entries
  4. Transactional errors did not include `fullUrl`

### Plugin and Facade API

1. Improved the message you get when the sorting/shaping operator is not implemented by your facade
2. `VonkOutcome` (and `VonkIssue`) has been simplified
3. `VonkConstants` has moved from `Vonk.Core.Context` to `Vonk.Core.Common`
4. `IResourceChangeRepository.Delete` requires a new second parameter: `string informationModel`. Information model constants can be found in `Vonk.Core.Common.VonkConstants.Model`
5. Exclude `Vonk.Fhir.R3` or `Vonk.Fhir.R4` from the `PipelineOptions` if you don't support it in your Facade.
6. Updated the minimal `PipelineOptions` for a Facade Plugin in the [example appsettings.json](#):
  - updated `Vonk.Core.Operations.SearchConfiguration` to `Vonk.Core.Operations.Search`
  - removed `Vonk.UI.Demo`
  - removed `Vonk.Core.Operations.Validate.SpecificationZipSourceConfiguration` from the `Exclude`
  - updated `Vonk.Core.Operations.Terminology` to `Vonk.Plugins.Terminology`

**Note:** Early Facade implementations were built with by using Vonk services and middleware in a self-built ASP.NET Core web server. This can be seen in the Vonk.Facade.Starter project in the [repository](#) with the same name. Due to changes in Vonk this does not work with Vonk 3.0.0. It will be fixed in 3.1.0. But after that such projects cannot be upgraded anymore and will have to be refactored to a proper plugin (as the ViSi.Repository project in the same repository). Please [contact](#) us in case of any questions.

### 3.2.17 Release 3.0.0-beta2

**Attention:** We updated the *Security notifications for Firely Server*.

#### Database

Note the changes in *3.0.0-beta1*, but there are no new changes in beta2.

#### Feature

1. *Subscriptions* works for R4 also. Note that a Subscription will only be activated for resource changes in the same FHIR version as the Subscription itself.
2. *Load Conformance Resources from disk* works for R4 also. Use a directory name that ends with .R4 for R4 conformance resources.
3. *Re-indexing for new or changed SearchParameters* works for R4 also. Issue a reindex with a fhirVersion parameter in the Accept header, and it will be executed for the SearchParameters defined for that FHIR version.
4. Allow for non-hl7 prefixed canonical urls for conformance resources (since sdf-7 is lifted). See *Custom Resources*.
5. Custom Resources can be validated, both individually and as part of a bundle. See *Custom Resources*.
6. If the Accept header lacks a *fhirVersion parameter*, it will fall back to the fhirVersion parameter of the Content-Type header and vice versa. If both are missing, Vonk will default to STU3.

#### Fix

1. `_include` did not work for R4.
2. `_include` gave a 500 response code if a resource contains absolute references.
3. A resource with unknown elements could result in an uncaught `Hl7.Fhir.ElementModel.StructuralTypeException`.
4. The homepage stated that Vonk was only for STU3. Fixed that.
5. Bundle.timestamp element (new in R4) was not populated in bundles returned from Search and History operations.
6. Some operations could return an `OperationOutcome` with an issue *and* a success message.
7. Better error message if a resource without any meta.profile is not accepted by *Validating incoming resources*.
8. Requesting an invalid FHIR version resulted in a `ArgumentNullException`.

### Plugin and Facade API

1. NuGet package `Vonk.Fhir.R4` had a dependency on `Vonk.Administration.API`, but the latter is not published. We removed the dependency.
2. `IResourceExtensions.UpdateMetadata` did not update the id of the resource.
3. `VonkOutcome.RemoveIssue()` method has been removed.

### Examples

1. Plugin example (`Vonk.Plugin.ExampleOperation`):
  1. Added an example of middleware directly interacting with the `HttpContext` (instead of just the `VonkContext`), see the file `VonkPluginMiddleware.cs`
  2. `CapabilityStatementBuilder` was not called.
2. DocumentOperation (`Vonk.Plugin.DocumentOperation`):
  1. Composition ID was not determined correctly when using POST.

### 3.2.18 Release 3.0.0-beta1

Vonk 3.0.0 is a major upgrade that incorporates handling FHIR R4. This runs in the same server core as FHIR STU3. See [Multiple versions of FHIR](#) for background info.

**Attention:** If you have overridden the `PipelineOptions` in your own settings, you should review the new additions to it in the `appsettings.default.json`. In particular we added `Vonk.Fhir.R4` that is needed to support FHIR R4.

**Attention:** MacOS: you may need to clean your temp folder from previous `specification.zip` expansions. Find the location of the temp folder by running `echo $TMPDIR`.

### Database

1. SQL Server, SQLite:
  1. `vonk.entry` got a new column 'InformationModel', set to 'Fhir3.0' for existing resources.
  2. `vonk.ref` got a new column 'Version'.
  3. Database indexes have been updated accordingly.

Vonk will automatically update both the Administration and the Data databases when you run Vonk 3.0.0.
2. MongoDB / CosmosDb:
  1. The documents in the `vonkentries` collection got a new element `im` (for `InformationModel`), set to 'Fhir3.0' for existing resources.
  2. The documents in the `vonkentries` collection got a new element `ref.ver` (for `Version`).
  3. Database indexes have been updated accordingly.
3. MongoDB / CosmosDb: Got a light mechanism of applying changes to the document structure. A single document is added to the collection for that, containing `VonkVersion` and `LatestMigration`.



4. MongoDB: The default name for the main database was changed from 'vonkstu3' to 'vonkdata'. If you want to continue using an existing 'vonkstu3' database, override `MongoDbOption:DatabaseName`, see [Hierarchy of settings](#).

## Feature

1. Support for FHIR R4 next to FHIR STU3. Vonk will choose the correct handling based on the `fhirVersion` parameter in the `mimetype`. The `mimetype` is read from the `Accept` header and (for POST/PUT) the `Content-Type` header. See [Multiple versions of FHIR](#) for background info.
2. Upgrade to HL7.Fhir.Net API 1.3, see its [releasenotes](#).
3. Administration API imports both STU3 and R4 conformance resources, see [Managing Conformance Resources](#)
  1. Note: [Terminology operations](#) are still only available for STU3.
  2. Note: [Subscriptions](#) are still only available for STU3.
4. Conditional delete on the Administration API. It works just as on the root, see [Create, read, update, patch, delete](#).
5. Defining a custom `SearchParameter` on a [Custom ResourceType](#) is now possible.
6. Canonical uris are now recognized when searching on references ([specification](#))
7. Vonk calls `UseIISIntegration` for better integration with IIS (if present).

## Fix

1. In the settings, `PipelineOptions.Branch.Path` for the root had to be `/`. Now you can choose your own base (like e.g. `/fhir`)
2. `$meta`:
  1. enabled on history endpoint (e.g. `/Patient/123/_history/v1`)
  2. disabled on type and system level
  3. returned empty `Parameters` resource if resource had no `meta.profile`, now returns the resources `meta` element.
  4. when called on a non-existing resource, returns 404 (was: empty `Parameters` resource)
  5. added to the `CapabilityStatement`
3. History on non-existing resource returned `OperationOutcome` instead of 404.
4. The setting for `SupportedInteractions` was not enforced for custom operations.
5. `CapabilityStatement.name` is updated from Vonk beta conformance to Vonk FHIR Server <version> `CapabilityStatement`.
6. [Terminology services](#):
  1. `$lookup` did not work on `GET /CodeSystem`
  2. `$lookup` did not support the `coding` parameter
  3. `$expand` did not fill in the expansion element.
  4. Operations were not listed in the `CapabilityStatement`.
  5. Namespace changed to `Vonk.Plugins.Terminology`, and adjusted accordingly in the default `PipelineOptions`.
7. A `SearchParameter` of type `token` did not work on an element of type `string`, e.g. `CodeSystem.version`.

8. Search with POST was broken.
9. If a long running task is active (response code 423, see *Import of Conformance Resources* and *Re-indexing for new or changed SearchParameters*), the OperationOutcome reporting that will now hide issues stating that all the arguments were not supported (since that is not the cause of the error).
10. Overriding an array in the settings was hard - it would still inherit part of the base setting if the base array was longer. We changed this: an array will always overwrite the complete base array. Note that this may trick you if you currently override a single array element with an environment variable. See *Hierarchy of settings*.
11. The element `meta.source` cannot be changed on subsequent updates to a resource (R4 specific)
12. `SearchParameter.StructureDefinition.ext-context` yielded many errors in the log because the definition of the fhirpath in the specification is not correct. We provided a corrected version in `errataFhir40.zip` (see *Errata to the specification*).
13. *Enable or disable interactions* was not evaluated for custom operations.
14. Delete of an instance accepted search parameters on the url.
15. Transactions: references to other resources in the transaction were not updated if the resource being referenced was in the transaction as an update (PUT). (this error was introduced in 2.0.0).

### Plugin and Facade API

1. A new NuGet package is introduced: `Vonk.Fhir.R4`.
2. `VonkConstants` moved to the namespace `Vonk.Core.Common` (was: `Vonk.Core.Context`)
3. `IResource.Navigator` element is removed (was already obsolete). Instead: Turn it into an `ITypedElement` and use that for navigation with `FhirPath`.
4. `InformationModel` element is added to
  1. `IResource`: the model in which the resource is defined (`VonkConstants.Model.FhirR3` or `VonkConstants.Model.FhirR4`)
  2. `IVonkContext`: the model that was specified in the Accept header
  3. `IModelService`: the model for which this service is valid (implementations are available for R3 and R4)
  4. `InteractionHandler` attribute: to allow you to specify that an operation is only valid for a specific FHIR version. This can also be done in the fluent interface with the new method `AndInformationModel`. See *Interaction Handling*
5. Dependency injection: if there are implementations of an interface for R3 and R4, the dependency injection in Vonk will automatically inject the correct one based on the `InformationModel` in the request.
6. If you want to register your own service just for one `informationmodel`, do that as follows:

Add a `ContextAware` attribute to the implementation class:

```
[ContextAware (InformationModels = new[] {VonkConstants.Model.FhirR3})]  
public class MySearchRepository{...}
```

Then register the service as being `ContextAware`:

```
services.TryAddContextAware<ISearchRepository, MySearchRepository>(ServiceLifeTime.  
    ↪ Scoped);
```

7. `FhirPropertyIndexBuilder` is moved to `Vonk.Fhir.R3` (and was already marked obsolete - avoid using it)

8. Implementations of the following that are heavily dependent upon version specific Hl7.Fhir libraries have been implemented in both Vonk.Fhir.R3 and Vonk.Fhir.R4.
  1. `IModelService`
  2. `IStructureDefinitionSummaryProvider` (to add type information to an `IResource` and turn it into an `ITypedElement`)
  3. `ValidationService`
9. `IConformanceContributor` is changed to `ICapabilityStatementContributor`. The methods on it have changed slightly as well because internally they now work on a version-independent model. Please review your `IConformanceContributor` implementations.

## Examples

1. Document plugin:
  1. *Document Bundle does not contain an identifier*
  2. *Missing unit test for custom resources*
  3. Upgraded to Vonk 2.0.0 libraries (no, not yet 3.0.0-beta1)
2. Facade example
  1. Added support for searching directly on a reference from Observation to Patient (e.g. `/Observation?patient=Patient/3`).
  2. Fixed support for `_revinclude` of Observation on Patient (e.g. `/Patient?_revinclude:Observation:subject:Patient`).
  3. Upgraded to Vonk 2.0.0 libraries (no, not yet 3.0.0-beta1)
3. Plugin example
  1. Added examples for pre- and post handlers.

## Known to-dos

1. *Re-indexing for new or changed SearchParameters*: does not work for R4 yet.
2. *Preloading resources*: does not work for R4 yet.
3. *Subscriptions*: do not work for R4 yet.
4. *Terminology services*: operations do not work for R4.
5. During *Import of Conformance Resources*: Files in the import directory and Simplifier projects are only imported for R3.

## 3.3 Old Firely Server release notes (v2.x and earlier)

### 3.3.1 Release 2.2.0

#### Database

1. SQL Server: the index tables have their clustered index on their link to the resources. SQL script '20190919000000\_Cluster\_Indexes\_On\_EntryId.sql' (found in the /data folder of the Vonk distribution) must be applied to existing Vonk SQL databases (both to the admin and to the data repositories) to make this change.

**Attention:** Vonk 2.2.0 (using SQL server) will not start unless this script has been applied to the databases. Please note that running the script can take considerable time, especially for large databases.

#### Features

1. When running Vonk in IIS, it now profits from the [in-process hosting model](#) that ASP.NET Core offers.
2. Further improved concurrent throughput for SQL Server.

#### Fix

1. On errors in a transaction, Vonk would not point out the entry if it had no fullUrl. Improved this message by using the resourcetype and id (if present) instead.
2. \_include gave a 500 responsecode if a resource contains absolute references.

### 3.3.2 Release 2.1.0

#### Database

1. SQL Server: Improved concurrent throughput.

#### Features

1. Upgrade to HL7.Fhir.Net API 1.3, see the [Older SDK release notes](#).
2. Vonk calls `UseIISIntegration` for better integration with IIS (if present).

#### Fix

1. Transactions: references to other resources in the transaction were not updated if the resource being referenced was in the transaction as an update (PUT). (this error was introduced in 2.0.0).

### 3.3.3 Release 2.0.1 hotfix

#### Fix

1. Supported Interactions were not checked for custom operations. In the *appsettings.json* the custom operations, like \$meta, were ignored. This has been fixed now.

### 3.3.4 Release 2.0.0 final

This is the final release of version 2.0.0, so the -beta is off. If you directly upgrade from version 1.1, please also review all the 2.0.0-beta and -beta2 release notes below.

**Attention:** We upgraded the version of .NET Core used to 2.2. Please get the latest 2.2.x runtime from the [.NET download site](#). The update was needed for several security patches and speed improvements.

**Attention:** The structure of the Validation section in the settings has changed. See [Validating incoming resources](#) for details.

**Attention:** This version of Vonk is upgraded to the Hl7.Fhir.API version 1.2.0. Plugin- and Facade builders will transitively get this dependency through the Vonk.Core package.

#### Database

No changes have been made to any of the database implementations.

#### Fix

1. When you created a StructureDefinition for a new resourcetype on /administration, the corresponding endpoint was not enabled.
2. Vonk does not update references in a transaction when a conditional create is used for an existing resource.
3. Paths in PipelineOptions would interfere if one was the prefix of the other.
4. Indexing a HumanName with no values but just extensions failed.
5. The selflink in a bundle did not contain the sort parameters. In this version the selflink always contains a sort and a count parameter, even if they were not in the request and the default values have been applied.
6. The import of conformance resources from specification.zip yielded warnings on .sch files.
7. Errors introduced in the 2.0.0-beta versions:
  1. Syntax errors in the XML or JSON payload yielded an exception, now they are reported with an OperationOutcome upon parsing.
  2. \$expand and other terminology operations caused a NullReference exception.
  3. \_element did not include the mandatory elements.

### Feature

1. Vonk supports Custom Resources. See [Custom Resources](#).
2. Operation *\$meta* is now supported, to quickly get the tags, security labels and profiles of a resource.
3. /metadata, retrieving the CapabilityStatement performs a lot better (just the initial call for a specific Accept-Type takes a bit longer).
4. Validation can be controlled more detailed. Choose the strictness of parsing independent of the level of validation. With this, the settings section 'Validation' has also changed. See [Validating incoming resources](#).

### Plugin and Facade API

1. We upgraded the embedded Fhir.Net API to version 1.2, see the [Older SDK release notes](#).
2. Together with the upgrade to .NET Core 2.2, several libraries were updated as well. Most notably Microsoft.EntityFrameworkCore.\*, to 2.2.3.

### 3.3.5 Release 2.0.0-beta2

#### Fix

- Fixed RelationalQuery in Vonk.Facade.Relational, so Vonk.Facade.Starter can be used again.

### 3.3.6 Release 2.0.0-beta

We have refactored Vonk internally to accomodate future changes. There are only minor functional changes to the FHIR Server. Facade and Plugin builders must be aware of a few interface changes, most notably to the IResource interface.

This release is a *beta* release because of the many internal changes, and because we expect to include a few more in the final release. Have a go with it in your test environment to see whether you encounter any trouble. We also encourage you to build your plugin and/or facade against it to prepare for code changes upon the final release.

You can still access the latest final release (1.1.0):

- Binaries: through the [Simplifier downloads page](#), choose 'List previous versions'.
- Docker: `docker pull simplifier/vonk:1.1.0`
- NuGet: `<PackageReference Include="Vonk.Core" Version="1.1.0" />`

#### Database

No changes have been made to any of the database implementations.

## Fix

1. The *\$validate* operation processes the profile parameter.
2. If an update brings a resource ‘back to life’, Vonk returns statuscode 201 (previously it returned 200).
3. On an initial Administration Import of specification.zip, Vonk found an error in valueset.xml. This file was fixed in the specification.zip that comes with Fhir.NET API 1.1.2.
4. Transaction: references within the transaction are automatically changed to the id’s the referenced resources get from Vonk when processing the transaction. This did not happen for references inside extensions. It does now.
5. Administration Import: an Internal Server Error could be triggered with a zip file with nested directories in it.
  - NB: Directories in your zip are still not supported because of [Fhir.NET API issue #883](#), but Vonk will not error on it anymore.
6. Search: The entry.fullUrl for an OperationOutcome in a Search bundle had a relative url.
7. Search: Processed \_elements and \_summary arguments were not reported in the selflink of the bundle (or any of the paging links).
8. Search: The selflink will include a \_count parameter, even if it was not part of the request and hence the default value for \_count from the *BundleOptions* was applied.
9. Search on :exact with an escaped comma (e.g. /Patient?name:exact=value1\,value2) was executed as a choice. Now the escape is recognized, and the argument processed as one term.

## Feature

1. Upgraded Fhir.NET API to version 1.1.2, see the [Older SDK release notes](#).
2. The Vonk Administration API now allows for StructureMap and GraphDefinition resources to be loaded.
3. The opening page of Vonk (and the only UI part of it) is updated. It no longer contains links that you can only execute with Postman, and it has a button that shows you the CapabilityStatement.
4. We published our custom operations on [Simplifier!](#) And integrated those links into the CapabilityStatement.
5. You can now access older versions of the Vonk binaries through the Simplifier downloads. (This was already possible for the Docker images and NuGet packages through their respective hubs).
6. [Vonk.IdentityServer.Test](#) and [Vonk.Facade.Starter](#) have been integrated into the Continuous Integration system.
7. In JSON, the order of the output has changed:
  1. If id and/or meta elements were added by Vonk (on a create or update), they will appear at the end of the resource.

## Plugin and Facade API

1. IResource interface and related classes have had several changes. If you encounter problems with adapting your code, please contact us.
  - It derives from the ISourceNode interface from the Fhir.NET API.
  - Change and Currency are properties that were only relevant in the repository domain, and not in the rest of the pipeline. They have been deprecated. You can access the values still with resource.GetChangeIndicator() and resource.GetCurrencyIndicator(). This is implemented with Annotations on the ISourceNode. All of Vonk’s own implementations retain those annotations, but if the relevant annotation is somehow missing, default values are returned (ResourceChange.NotSet resp. ResourceCurrency.Current).

- The Navigator property is obsolete. The type of it (IElementNavigator) is obsolete in the Fhir.NET API. To run FhirPath you provide type information and run the FhirPath over an ITypedElement:

```
//Have IStructureDefinitionSummaryProvider _schemaProvider injected in the
↪ constructor.
var typed = resource.ToTypedElement(_schemaProvider);
var matchingElements = typed.Select('your-fhirpath-expression');
```

- Id, Version and LastUpdated can no longer be set directly on the IResource instance. IResource has become **immutable** (just like ISourceNode). The alternatives are:

```
var resourceWithNewId = resource.SetId("newId");
var resourceWithNewVersion = resource.SetVersion("newVersion");
var resourceWithNewLastUpdated = resource.SetLastUpdated(DateTimeOffset.UtcNow);
```

- Because the IChangeRepository is responsible for creating new id's and versions, we also included extensions methods on it to update all three fields at once:

```
var updatedeResource = changeRepository.EnsureMeta(resource, KeepExisting.Id /
↪ Version / LastUpdated);
var updatedResource = changeRepository.FreshMeta(resource); //replaces all three
```

2. The PocoResource class is obsolete. To go from a POCO (like an instance of the Patient class) to an IResource, use the ToIResource() extension method found in Vonk.Fhir.R3.
3. The PocoResourceVisitor class is obsolete. Visiting can more effectively be done on an ITypedElement:

```
//Have IStructureDefinitionSummaryProvider _schemaProvider injected in the
↪ constructor.
var typed = resource.ToTypedElement(_schemaProvider);
typed.Visit((depth, element) => { //do what you want with element});
```

4. SearchOptions has changed:
  - Properties Count and Offset have been removed.
  - Instead, use \_count and \_skip arguments in the IArgumentCollection provided to the SearchRepository.Search method if you need to.
5. We have created a template for a plugin on [GitHub](#). Fetch it for a quick start of your plugin.

### 3.3.7 Release 1.1.0

**Attention:** New security issues have been identified by Microsoft. See the [Security notifications for Firely Server](#) for details.

**Attention:** The setting for the location of the license file has moved. It was in the top level setting LicenseFile. It still has the same name, but it has moved into the section License. See [License](#) for details.

**Attention:** This version of Vonk is upgraded to the Hl7.Fhir.API version 1.1.1. Plugin- and Facade builders will transitively get this dependency through the Vonk.Core package.



## Database

No changes have been made to any of the database implementations.

## Feature

1. Vonk will count the number of requests that it processes. See [License](#) for settings on that. Because of this change, the `LicenseFile` setting has moved from the top level to under `License`.
2. The plugin folder ([Configuring the Firely Server Pipeline](#)) may now contain subfolders. Plugins will be read from all underlying folders.
3. Vonk supports If-Match on update. See [Managing Resource Contention](#) in the specification for details.
4. Plugins may return non-FHIR content. See [Returning non-FHIR content from a plugin](#).
5. This feature may also be used for [Custom Authentication](#).
6. A [Template for a plugin](#) is added to the documentation.
7. A documentation page on performance is added.
8. Upgrade of the `HL7.Fhir.API` library to 1.1. See the [Older SDK release notes](#).

## Fix

1. Transaction: forward references from one resource to another in a Transaction were not correctly resolved.
2. When you set `ValidateIncomingResources` to true, Vonk no longer accepts resources with extensions that are unknown to it. This is now also reflected in the `CapabilityStatement.acceptUnknown`.
3. The links in a bundle response (`Bundle.link`) were relative links. Now they are absolute links.
4. HTTP 500 instead of an OO was returned when trying to update a subscription with an invalid request status.
5. If an error is found in a `SearchParameter` in the Administration database, Vonk logs the (canonical) url of that `SearchParameter` for easier reference.
6. Transaction: Response bundle contained versioned fullUrls. We changed that to unversioned urls.
7. Bundles: Response bundles with an `OperationOutcome` contained a versioned fullUrl for the entry containing the `OperationOutcome`. We changed that to an unversioned url.
8. Deleting a resource from the Administration API that does not exist would lead to an internal server error.

## Supported Plugins

1. Several fixes have been done on the [Document plugin](#).

### 3.3.8 Release 1.0.0

Yes! Vonk version 1.0 is out. It is also the first version that is released without the -beta postfix. It has been very stable from the very first version, and now we think it is time to make that formal.

Release 1.0.0 is functionally identical to 0.7.4.0. But we optimized the deployment process for *Firely Server for your Simplifier project* and *Docker* in general. The contents of the core specification are now preloaded in the SQLite administration database, so your first startup experience is a lot faster.

### 3.3.9 Release 0.7.4.0

#### Database

1. The index definitions for SQL Server have been updated for improved performance. This should be handled automatically when you start Vonk 0.7.4 and have *AutoUpdateDatabase* enabled.

#### Fix

1. Posting a resource with an invalid content-type to the regular FHIR endpoint should result in HTTP 415 and not HTTP 400.
2. Warning ‘End method “PocoResourceVisitor.VisitByType”, could not cast entity to PocoResource.’ in the log was incorrect.
3. When running Administration API on SQLite and Vonk on SQL Server, update or delete would fail.
4. Handle quantity with very low precision (e.g. ‘3 times per year’ - 3|http://unitsofmeasure.org/a).
5. POST to <vonk\_base>/Administration/\* with another Content-Type than application/json or application/xml results in HTTP 500.

#### Feature

1. Support forward references in a *Transaction bundle*. Previously Vonk would only process references back to resources higher up in the bundle.
2. Performance of Validation and Snapshot Generation has improved by approximately 10 times...
3. ... and correctness has improved as well.
4. Administration API also support the NamingSystem resource.

### 3.3.10 Release 0.7.3.0

#### Fix

1. Search on /administration/Subscription was broken
2. Neater termination of the Subscription evaluation process upon Vonk shutdown
3. A Bundle of type batch is now rejected if it contains internal references.
4. Urls in the narrative (href and src) are also updated to the actual location on the server.
5. A system wide search on compartment returns 403, explaining that that is too costly.

### 3.3.11 Release 0.7.2.1

#### Fix

1. Delete on /administration was broken.

### 3.3.12 Release 0.7.2.0

#### Database

1. Fixes 2 and 3 require a reindex for specific searchparameters, if these parameters are relevant to you.

#### Features and fixes

1. Fix: Reject a search containing a modifier that is incorrect or not supported.
2. Fix: The definition for searchparameter Encounter.length was unclear. We added the correct definition from FHIR R4 to the errata.zip, so it works for STU3 as well. If this is relevant for you, you may want to reindex for this searchparameter. See [Rebuild the search index for specific searchparameters](#), just for 'Encounter.length'.
3. Fix: Error "Unable to index for element of type 'base64Binary'". This type of element is now correctly indexed. One known searchparameter that encounters this type is Device.udi-carrier. If this is relevant to you, you may want to reindex for this searchparameter. See [Rebuild the search index for specific searchparameters](#), just for 'Device.udi-carrier'.
4. Fix: Validation would fail on references between contained resources. See also fix #423 in the [Older SDK release notes](#).
5. Fix: E-tag was missing from the response on a delete interaction.
6. Fix: An invalid mimetype in the \_format parameter (like \_format=application/foobar) returned response code 400 instead of 415.
7. Fix: If a subscription errors upon execution, not only set the status to error, but also state the reason in Subscription.error for the user to inspect.
8. Fix: Search on /Observation?value-string:missing=false did not work. As did the missing modifier on other searchparameters on value[x] elements.
9. Feature: After /administration/importResources (see [Load Conformance Resources on demand](#)), return an OperationOutcome detailing the results of the operation.
10. Feature: Upon usage of a wrong value for \_summary, state the possible, correct values in the OperationOutcome.
11. Feature: Allow for multiple deletes with a Conditional Delete, see [Create, read, update, patch, delete](#).
12. Feature: The version of Vonk is included in the log file, at startup.
13. Configuration: Add Vonk.Smart to the PipelineOptions by default, so the user only needs to set the SmartAuthorizationOptions.Enabled to true.
14. Upgrade: We upgraded to the latest C# driver for MongoDB (from 2.4.4 to 2.7.0).

### 3.3.13 Release 0.7.1.1

#### Fix

Spinning up a Docker container would crash the container because there was no data directory for SQLite (the default repository). This has been solved now: Vonk will create the data directory when it does not exist.

### 3.3.14 Release 0.7.1.0

**Attention:** Fix nr. 8 requires a `reindex/searchparameters` with `include=Resource._id,Resource._lastUpdated,Resource._tag`. Please review *Re-indexing for new or changed SearchParameters* on how to perform a reindex and the cautions that go with it. Also note the changes to reindexing in fix nr. 1.

#### Database

1. We added support for SQLite! See *Using SQLite* for details.
2. We also made SQLite the default setting for both the main Vonk database and the *Firely Server Administration API*.
3. With the introduction of SQLite we advise running the Administration API on SQLite. In the future we will probably deprecate running the Administration API on any of the other databases.
4. Support for CosmosDB is expanded, though there are a few limitations.

#### Facade

1. If you rejected the value for the `_id` searchparameter in your repository, Vonk would report an `InternalServerError`. Now it reports the actual message of your `ArgumentException`.

#### Features and fixes

1. We sped up *Re-indexing for new or changed SearchParameters*. The request will be responded to immediately, while Vonk starts the actual reindex asynchronously and with many threads in parallel. Users are guarded against unreliable results by blocking other requests for the duration of the reindex. Reindexing is still not to be taken lightly. It is a **very heavy** operation that may take very long to complete. See *Re-indexing for new or changed SearchParameters* for details.
2. A really large bundle could lead Vonk (or more specifically: the validator in Vonk) to a `StackOverflow`. You can now set *limits* to the size of incoming data to avoid this.
3. *Reindexing* is supported on CosmosDB, but it is less optimized than on MongoDB.
4. Using `_include` or `_revinclude` would yield an `OperationOutcome` if there are no search results to include anything on. Fixed that to return 404 as it should.
5. Using the `:not` modifier could return false positives.
6. A batch or transaction with an entry having a value for `IfModifiedSince` would fail.
7. History could not be retrieved for a deleted resource. Now it can.
8. *Reindex* would ignore the generic searchparameters defined on `Resource` (`_id`, `_lastUpdated`, `_tag`). Because `id` and `lastUpdated` are also stored apart from the search index, this was really only a problem for `_tag`. If you rely on the `_tag` searchparameter you need to reindex **just for the searchparameter ``Resource.\_tag``**.

9. Vonk logs its configuration at startup. See [Log of your configuration](#) for details.

### 3.3.15 Release 0.7.0.0

#### Database

1. Indexes on the SQL Server repository were updated to improve performance. They will automatically be applied with [AutoUpdateDatabase](#).

#### Facade

1. Release 0.7.0.0 is compatible again with Facade solutions built on the packages with versions 0.6.2, with a few minor changes. Please review the Vonk.Facade.Starter project for an example of the necessary adjustments. All the differences can be seen in [this file comparison](#).
2. Fix: The SMART authorization failed when you don't support all the resourcetypes. It will now take into account the limited set of supported resourcetypes.
3. Fix: Vonk.Facade.Relational.RelationalQueryFactory would lose a `_count` argument.
4. Documentation: We added documentation on how to implement Create, Update and Delete in a facade on a relational database. See [Enable changes to the repository](#). This is also added to the [example Facade solution](#) on GitHub.

#### Features and fixes

1. Feature: [Vonk FHIR Plugins](#) has been released. You can now add libraries with your own plugins through configuration.
2. Feature: Through [Vonk FHIR Pluginss](#) you can replace the landing page with one in your own style. We provided an [example](#) on how to do that.
3. Feature: You can now start Vonk from within another directory than the Vonk binaries directory, e.g. `c:\programs>dotnet .\vonk\vonk.server.dll`.
4. Feature: You can configure the maximum number of entries allowed in a Batch or Transaction, to avoid overloading Vonk. See [Batch and transaction](#).
5. Upgrade: We upgraded the FHIR .NET API to version 0.96.0, see the [Older SDK release notes](#) for details. Mainly #599 affects Vonk, since it provides the next...
6. Fix: Under very high load the FhirPath engine would have concurrency errors. The FhirPath engine is used to extract the search parameters from the resources. This has been fixed.
7. Fix: Search on a frequently used tag took far too long on a SQL Server repository.
8. Fix: The [Patient.deceased](#) search parameter from the specification had an error in its FhirPath expression. We put a corrected version in the [errata.zip](#).
9. Fix: Several composite search parameters on Observation are defined incorrectly in the specification, as is reported in [GForge issue #16001](#). Until the specification itself is corrected, we provide corrections in the [errata.zip](#).
10. Fix: Relative references in a resource that start with a forward slash (like `/Patient/123`) could not be searched on.
11. Fix: System wide search within a compartment looked for the pattern `<base>/Patient/123/?_tag=bla`. Corrected this to `<base>/Patient/123/*?_tag=bla`

12. Fix: When loading *Simplifier resources*, Vonk can now limit this to the changes since the previous import, because the Simplifier FHIR endpoint supports `_lastUpdated`.
13. Fix: *Conformance resources* are always loaded into the Administration API when running on a Memory repository. Or actually, always if there are no StructureDefinitions in the Administration database. To enable this change, imported files are no longer moved to the *AdministrationOptions.ImportedDirectory*.
14. Fix: *Re-indexing for new or changed SearchParameters* would stop if a resource was encountered that could not properly be indexed. It will now continue working and report any errors afterwards in an *OperationOutcome*.
15. Fix: The terms and privacy statement on the default landing page have been updated.
16. Fix: When searching on a search parameter of type date, with an argument precision to the minute (but not seconds), Vonk would reject the argument. It is now accepted.
17. Fix: DateTime fields are always normalized to UTC before they are stored. This was already the case on MongoDB, and we harmonized SQL and Memory to do the same. There is no need to reindex for this change.
18. Fix: When you use accents or Chinese characters in the url for a search, Vonk gives an error.
19. Fix: A reverse chained search on MongoDB sometimes failed with an Internal Server Error.

### 3.3.16 Release 0.6.5.0

**Attention:** This version changes the way conformance resources are loaded from zip files and/or directories at startup. They are no longer loaded only in memory, but are added to the Administration API's database. You will notice a delay at first startup, when Vonk is loading these resources into the database. See Feature #1 below.

**Attention:** 2018-06-07: We updated the Database actions for 0.6.5.0, you should always perform a reindex, see right below.

#### Database

1. Feature 2, 4 and 14 below require a *reindex/all*, both for MongoDB and SQL Server.

#### Facade

1. Release 0.6.5.0 is not released on NuGet, so the latest NuGet packages have version 0.6.2-beta. Keep an eye on it for the next release...

#### Features and fixes

1. Feature: Run Vonk from you Simplifier project! See *Use Firely Server with your Simplifier artifacts* for details.
2. Feature: Vonk supports Microsoft Azure CosmosDB. This required a few small changes to the MongoDB implementation (the share the drivers), so please reindex your MongoDB database: *reindex/all*.
3. Feature: Configuration to restrict support for ResourceTypes, SearchParameters and CompartmentDefinitions, see *Restrict supported resources and SearchParameters*.
4. Feature: Errata.zip: collection of corrected search parameters (e.g. that had a faulty expression in the FHIR Core specification), see *Errata to the specification*

5. Upgrade: FHIR .NET API 0.95.0 (see the [Older SDK release notes](#))
6. Fix: a search on `_id:missing=true` was not processed correctly.
7. Fix: better distinction of reasons to reject updates (error codes 400 vs. 422, see [RESTful API specification](#))
8. Fix: recognize `_format=text/xml` and return xml (instead of the default json)
9. Fix: handling of the `:not` modifier in token searches (include resource that don't have a value at all).
10. Fix: handling of the `:not` modifier in searches with choice arguments
11. Fix: `fullUrl` in return bundles cannot be version specific.
12. Fix: evaluate `_count=0` correctly (it was ignored).
13. Fix: correct error message on an invalid `_include` (now Vonk tells you which resourcetypes are considered for evaluating the used searchparameter).
14. Fix: indexing of `Observation.combo-value-quantity` failed for UCUM code for Celcius. This fix requires a [reindex/all](#) on this searchparameter.
15. Fix: total count in history bundle.
16. Fix: on `vonk.fire.ly` we disabled validating all input, so you can now create or update resources also if the relevant profiles are not loaded (this was necessary for Crucible, since it references US Core profiles, that are not present by default).
17. Fix: timeout of Azure Web App on first startup of Vonk - Vonk's first startup takes some time due to import of the specification (see [Default Conformance Resources](#)). Since Azure Web Apps are allowed a startup time of about 3 minutes, it failed if the web app was on a low level service plan. Vonk will now no longer await this import. It will finish startup quickly, but until the import is finished it will return a 423 'Locked' upon every request.
18. Fix: improved logging on the import of conformance resources at startup (see [Import of Conformance Resources](#)).

### 3.3.17 Release 0.6.4.0

**Attention:** This version changes the way conformance resources are loaded from zip files and/or directories at startup. They are no longer loaded only in memory, but are added to the Administration API's database. You will notice a delay at first startup, when Vonk is loading these resources into the database. See Feature #1 below.

#### Database

1. Fix #9 below requires a [reindex/all](#).

#### Facade

1. Release 0.6.4.0 is not released on NuGet, so the latest NuGet packages have version 0.6.2-beta. This release is targeted towards the Administration API and [Terminology services](#), both of which are not (yet) available in Facade implementations. We are working on making the features of the Administration API available to Facade implementers in an easy way.

### Features and fixes

1. Feature: Make all loaded conformance resources available through the Administration API.

Previously:

- Only SearchParameter and CompartmentDefinition resources could be loaded from ZIP files and directories;
- And those could not be read from the Administration API.

Now:

- The same set of (conformance) resourcetypes can be read from all sources (ZIP, directory, Simplifier);
- They are all loaded into the Administration database and can be read and updated through the Administration API.

Refer to *Managing Conformance Resources* for details.

2. Feature: Experimental support for *Terminology services* operations \$validate-code, \$expand, \$lookup, \$compose.
3. Feature: Support for *Compartment Search*.
4. Feature: Track timing of major dependencies in *Azure Application Insights*.
5. Feature: *Log settings* can be overridden in 4 levels, just as the appsettings. The logsettings.json file will not be overwritten anymore by a Vonk distribution.
6. Fix: The check for *allowed profiles* is no longer applied to the Administration API. Previously setting Allowed-Profiles to e.g. [<http://mycompany.org/fhir/StructureDefinition/mycompany-patient>] would prohibit you to actually create or update the related StructureDefinition in the Administration API.
7. Fix: When posting any other resourcetype than the supported conformance resources to the Administration API, Vonk now returns a 501 (Not Implemented).
8. Fix: Support search on Token with only a system (e.g. <base>/Observation?code=http://loinc.org|)
9. Fix: Support search on Token with a fixed system, e.g. <base>/Patient?gender=http://hl7.org/fhir/codesystem-administrative-gender.html|female. This fix requires a *reindex/all*.
10. Fix: Reindex could fail when a Reference Searchparameter has no targets.
11. Fix: Vonk works as Data Server on *ClinFHIR*, with help of David Hay.
12. Fix: Clearer error messages in the log on configuration errors.
13. Fix: Loading conformance resources from disk in Docker.

### Documentation

1. We added documentation on *using IIS or NGINX as reverse proxies* for Vonk.
2. We added documentation on running Vonk on Azure Web App Services.



### 3.3.18 Release 0.6.2.0

**Attention:** The loading of appsettings is more flexible. After installing a new version you can simply paste your previous appsettings.json in the Vonk directory. Vonk's default settings are now in appsettings.default.json. see [Firely Server settings](#) for details.

#### Database

No changes

#### Features and fixes

1. Feature: Conditional References in *Transactions* are resolved.
2. Feature: More flexible support for different serializers (preparing for ndjson in Bulkdata)
3. Feature: Improved handling on missing settings or errors in the [Firely Server settings](#).
4. Feature: Improved *logging*, including Dependency Tracking on Azure Application Insights, see [Application Insights](#)
5. Feature: SearchParameter and CompartmentDefinition are now also imported from *Simplifier*, so both Simplifier import and the *Administration API* support the same set of conformance resources: StructureDefinition, SearchParameter, CompartmentDefinition, ValueSet and CodeSystem. See [Conformance resources](#).
6. Feature: Loading of appsettings is more flexible, see [Firely Server settings](#).
7. Feature: Added documentation on running Vonk behind IIS or NGINX: [Deploy Firely Server on a reverse proxy](#).
8. Performance: Improvement in speed of validation, especially relevant if you are *Validating incoming resources*.
9. Fix: If you try to load a SearchParameter (see [Load Conformance Resources from disk](#)) that cannot be parsed correctly, Vonk puts an error about that in the log.
10. Fix: Results from \_include and \_revinclude are now marked with searchmode: Include (was incorrectly set to 'Match' before)
11. Fix: \_format as one of the parameters in a POST Search is correctly evaluated.
12. Fix: No more errors in the log about a Session being closed before the request has finished ("Error closing the session. System.OperationCanceledException: The operation was canceled.")
13. Fix: Subscription.status is evaluated correctly upon create or update on the Administration API
14. Fix: Token search with only a system is supported (Observation.code=somesystem|)
15. Fix: On validation errors like 'Cannot resolve reference Organization/Organization-example26'' are now suppressed since the validator is set not to follow these references.
16. Fix: New Firely logo in SVG format - looks better
17. Fix: Creating resources with duplicate canonical url's on the Administration API is prohibited, see [Managing Conformance Resources](#).
18. Fix: If a Compartment filter is used on a parameter that is not implemented, Vonk will return an error, see [Compartments](#).

### 3.3.19 Release 0.6.1.0

Name change from Furore to Firely

### 3.3.20 Release 0.6.0.0

#### Attention:

- SearchParametersImportOptions is renamed to *MetadataImportOptions*.
- *Subscription* can now be disabled from the settings.

#### Database

1. The MongoDB implementation got a new index. It will be created automatically upon startup.

#### Features and fixes

1. Feature: *Access control based on SMART on FHIR*.
2. Feature: Vonk can also load CompartmentDefinition resources. See *Managing Conformance Resources* for instructions.
3. Feature: ValueSet and CodeSystem resources can be loaded into the administration endpoint, and loaded from Simplifier. See *Managing Conformance Resources* for instructions.
4. Feature: Be lenient on trailing slashes in the url.
5. Feature: OperationOutcome is now at the top of a Bundle result. For human readers this is easier to spot any errors or warnings.
6. Fix: In the *settings for SQL Server* it was possible to specify the name of the Schema to use for the Vonk tables. That was actually not evaluated, so we removed the option for it. It is fixed to 'vonk'.
7. Fix: The OperationOutcome of the *Reset* operation could state both an error and overall success.
8. Fix: If you did not set the CertificatePassword in the appsettings, Vonk would report a warning even if the password was not needed.
9. Fix: *Loading conformance resources* in the SQL Server implementation could lead to an error.
10. Fix: Clearer error messages if the body of the request is mandatory but empty.
11. Fix: Clearer error message if the Content-Type is missing.
12. Fix: GET on [base]/ would return the UI regardless of the Accept header. Now if you specify a FHIR mimetype in the Accept header, it will return the result of a system wide search.
13. Fix: In rare circumstances a duplicate logical id could be created.
14. Fix: GET [base]/metadat would return status code 200 (OK). But it should return a 400 and an OperationOutcome stating that 'metadat' is not a supported resource type.

## Documentation

1. We consolidated documentation on loading conformance resources into *Managing Conformance Resources*.

### 3.3.21 Release 0.5.2.0

**Attention:** Configuration setting SearchOptions is renamed to BundleOptions.

## Features and fixes

1. Fix: When you specify LoadAtStartup in the *ResourceLoaderOptions*, an warning was displayed: “WRN No server base configured, skipping resource loading.”
2. Fix: *Conditional create* that matches an existing resource returned that resource instead of an OperationOutcome.
3. Fix: *\_has*, *\_type* and *\_count* were in the CapabilityStatement twice.
4. Fix: *\_elements* would affect the stored resource in the Memory implementation.
5. Fix: Getting a resource with an invalid id (with special characters or over 64 characters) now returns a 404 instead of 501.
6. Feature: *Re-indexing for new or changed SearchParameters* now also re-indexes the Administration API database.
7. Fix: modifier :above for parameter type Url now works on the MongoDB implementation.
8. Fix: Vonk would search through inaccessible directories for the specification.zip.
9. Fix: Subscription could not be posted if ‘Database’ was not one of the SearchParametersImportOptions.
10. Fix: *\_(rev)include=\** is not supported but was not reported as such.
11. Fix: In a searchresult bundle, the references to other resources are now made absolute, refering to the Vonk server itself.
12. Fix: *BundleOptions* (previously: SearchOptions) settings were not evaluated.
13. Fix: Different responses for invalid resources when you change ValidateIncomingResources setting (400 vs. 501)
14. Fix: Better reporting of errors when there are invalid modifiers in the search.
15. Fix: Creating a resource that would not fit MongoDB’s document size resulted in an inappropriate error.
16. Fix: There was no default sort order in the search, resulting in warnings from the SQL implementation. Added default sort on *\_lastUpdated* (desc).
17. Fix: Preliminary disposal of LocalTerminology server by the Validator.

### Facade

1. Fix: `_include/_revinclude` on searchresults having contained resources triggered a `NotSupportedException`.

### 3.3.22 Release 0.5.1.1

#### Facade

We released the Facade libraries on [NuGet](#) along with *getting started documentation*.

No features have been added to the Vonk FHIR Server.

### 3.3.23 Release 0.5.0.0

#### Database

1. Long URI's for token and uri types are now supported, but that required a change of the SQL Server database structure. If you have `AutoUpdateDatabase` enabled (see *Using SQL server*), Vonk will automatically apply the changes. As always, perform a backup first if you have production data in the database.
2. To prevent duplicate resources in the database we have provided a unique index on the Entry table. This update does include a migration. It can happen that that during updating of your database it cannot apply the unique index, because there are duplicate keys in your database (which is not good). Our advise is to empty your database first (with `<vonk-endpoint>/administration/reset`, then update Vonk with this new version and then run Vonk with `AutoUpdateDatabase=true` (for the normal and the administration databases).

If you run on production and encounter this problem, please contact our support.

#### Features and fixes

1. Feature: POST on `_search` is now supported
2. Fix: Statuscode of `<vonk-endpoint>/administration/preload` has changed when zero resources are added. The statuscode is now 200 instead of 201.
3. Fix: OPTIONS operation returns now the capability statement with statuscode 200.
4. Fix: A search operation with a wrong syntax will now respond with statuscode 400 and an `OperationOutcome`. For example GET `<vonk-endpoint>/Patient?birthdate<1974` will respond with statuscode 400.
5. Fix: A statuscode 501 could occur together with an `OperationOutcome` stating that the operation was successful. Not anymore.
6. Fix: An `OperationOutcome` stating success did not contain any issue element, which is not valid. Solved.
7. Improvement: In the configuration on *Load Conformance Resources from simplifier.net* the section `ArtifactResolutionOptions` has changed to `ResourceLoaderOptions` and a new option has been introduced under that section named `LoadAtStartup` which, if set to true, will attempt to load the specified resource sets when you start Vonk
8. Improvement: the Memory implementation now also supports `SimulateTransactions`
9. Improvement: the option `SimulateTransactions` in the configuration defaults to false now
10. Feature: You can now add `SearchParameters` at runtime by POSTing them to the Administration API. You need to apply *Re-indexing for new or changed SearchParameters* to evaluate them on existing resources.
11. Fix: The batch operation with search entries now detects the correct interaction.

12. Fix: ETag header is not sent anymore if it is not relevant.
13. Fix: Searching on a String SearchParameter in a MongoDB implementation could unexpectedly broaden to other string parameters.
14. Fix: If Reference.reference is empty in a Resource, it is no longer filled with Vonks base address.
15. Feature: Search operation now supports `_summary`.
16. Fix: Paging is enabled for the history interaction.
17. Fix: Conditional updates won't create duplicate resources anymore when performing this action in parallel.
18. Fix: Indexing of CodeableConcept has been enhanced.
19. Fix: Search on reference works now also for an absolute reference.
20. Fix: Long uri's (larger than are 128 characters) are now supported for Token and Uri SearchParameters.
21. Improvement: The configuration of IP addresses in *Limited access* has changed. The format is no longer a comma-separated string but a proper JSON array of strings.

### 3.3.24 Release 0.4.0.1

#### Database

1. Long URL's for absolute references are now supported, but that required a change of the SQL Server database structure. If you have AutoUpdateDatabase enabled, Vonk will automatically apply the changes. As always, perform a backup first if you have production data in the database.
2. Datetime elements have a new serialization format in MongoDB. After installing this version, you will see warnings about indexes on these fields. Please perform *Re-indexing for new or changed SearchParameters*, for all parameters with `<vonk-endpoint>/administration/reindex/all`. After the operation is complete, restart Vonk and the indexes will be created without errors.

#### Features and fixes

1. Fix: SearchParameters with a hyphen ('-', e.g. general-practitioner) were not recognized in (reverse) chains.
2. Fix: CapabilityStatement is more complete, including (rev)includes and support for generic parameters besides the SearchParameters (like `_count`). Also the SearchParameters now have their canonical url and a description.
3. Improvement: *Preloading resources* gives more informative warning messages.
4. Fix: *Re-indexing for new or changed SearchParameters* did not handle contained resources correctly. If you have used this feature on the 0.3.3 version, please apply it again with `<vonk-endpoint>/administration/reindex/all` to correct any errors.
5. Improvement: *Loading resources from Simplifier* now also works for the Memory implementation.
6. Improvements on *Validation*:
  - profile parameter can also be supplied on the url
  - if validation is successful, an OperationOutcome is still returned
  - it always returns 200, and not 422 if the resource could not be parsed
7. Feature: support for Conditional Read, honouring if-modified-since and if-none-match headers.
8. Fix: Allow for url's longer than 128 characters in Reference components.
9. Fix: Allow for an id in a resource on a Create interaction (and ignore that id).

10. Fix: Allow for an id in a resource on a Conditional Update interaction (and ignore that id).
11. Fix: Include Last-Modified header on Capability interaction.
12. Fix: Format Last-Modified header in `httpdate` format.
13. Fix: Include version in `bundle.entry.fullUrl` on the History interaction.
14. Fix: Update `_sort` syntax from DSTU2 to STU3. Note: `_sort` is still only implemented for `_lastUpdated`, mainly for the History interaction.
15. Improvement: If the request comes from a browser, the response is sent with a Content-Type of `application/xml`, to allow the browser to render it natively. Note that most browsers only render the narrative if they receive xml.

### 3.3.25 Release 0.3.3.0

**Attention:** We upgraded to .NET Core 2.0. For this release you have to install .NET Core Runtime 2.0, that you can download from [dot.net](https://dot.net).

## Hosting

The options for enabling and configuring HTTPS have moved. They are now in `appsettings.json`, under ‘Hosting’:

```
"Hosting": {  
  "HttpPort": 4080,  
  "HttpsPort": 4081, // Enable this to use https  
  "CertificateFile": "<your-certificate-file>.pfx", //Relevant when HttpsPort is  
  present  
  "CertificatePassword" : "<cert-pass>" // Relevant when HttpsPort is present  
},
```

This means you have to adjust your environment variables for `CertificateFile` and `CertificatePassword` (if you had set them) to:

`VONK_Hosting:CertificateFile`  
`VONK_Hosting:CertificatePassword`

The setting ‘UseHttps’ is gone, in favour of `Hosting:HttpsPort`.

## Database

There are no changes to the database structure.

## Features and fixes

1. Feature: Subscription is more heavily checked on create and update. If all checks pass, status is set to active. If not, the Subscription is not stored, and Vonk returns an OperationOutcome with the errors.
  - Criteria must all be supported
  - Endpoint must be absolute and a correct url
  - Enddate is in the future
  - Payload mimetype is supported
2. Feature: use `_elements` on Search
3. Feature: *load profiles from your Simplifier project* at startup.
4. Feature: Content-Length header is populated.
5. Fix: PUT or POST on `/metadata` returned 200 OK, but now returns 405 Method not allowed.
6. Fix: Sometimes an error message would appear twice in an OperationOutcome.
7. Fix: `_summary` is not yet implemented, but was not reported as ‘not supported’ in the OperationOutcome. Now it is. (Soon we will actually implement `_summary`.)
8. Fix: If-None-Exist header was also processed on an update, where it is only defined for a create.
9. Fix: Set `Bundle.entry.search.mode` to ‘outcome’ for an OperationOutcome in the search results.
10. UI: Display software version on homepage.

### 3.3.26 Release 0.3.2.0

1. Fix: `_include` and `_revinclude` could include too many resources.

### 3.3.27 Release 0.3.1.0

1. IP address restricted access to Administration API functions.
2. Fix on Subscriptions:
  1. Accept only Subscriptions with a channel of type `rest-hook` and the payload (if present) has a valid mime-type.
  2. Set them from requested to active if they are accepted.

### 3.3.28 Release 0.3.0.0

1. Database changes
 

If you have professional support, please consult us on the best way to upgrade your database.

  1. The schema for the SQL Database has changed. It also requires re-indexing all resources.
  2. The (implicit) schema for the documents in the MongoDB database has changed.
  3. The Administration API requires a separate database (SQL) or collection (MongoDb).
2. New features:
  1. *Custom Search Parameters*

2. Support for Subscriptions with rest-hook channel
3. Preload resources from a zip.
4. Reset database
5. Conditional create / update / delete
6. Support for the prefer header
7. Validation on update / create (can be turned on/off)
8. Restrict creates/updated to specific profiles.
9. Configure supported interactions (turn certain interactions on/off)
3. New search features:
  1. `_has`
  2. `_type` (search on system level)
  3. `_list`
  4. `_revinclude`
4. Enhancements
  1. `:exact`: Correctly search case (in)sensitive when the `:exact` modifier is (not) used on string parameters.
  2. Enhanced reporting of errors and warnings in the `OperationOutcome`.
  3. Custom profiles / `StructureDefinitions` separated in the Administration API (instead of in the regular database).
  4. Full FHIRPath support for Search Parameters.
  5. Fixed date searches on dates without seconds and timezone
  6. Fixed evaluation of modifier `:missing`
  7. Correct total number of results in search result bundle.
  8. Fix paging links in search result bundle
  9. Better support for mimetypes.
5. DevOps:
  1. New *Firely Server Administration API*
  2. Enabled logging of the SQL statements issued by Vonk (see *Log settings*)
  3. Migrations for SQL Server (auto create database schema, also for the Administration API)
6. Performance
  1. Added indexes to MongoDB and SQL Server implementations.



## 3.4 Security notifications for Firely Server

### 3.4.1 January 2021

Microsoft has a new Security Advisory regarding ASP.NET Core:

- Microsoft Security Advisory CVE-2020-1161 | ASP.NET Core Denial of Service Vulnerability in all ASP.NET Core applications on runtime 3.1.3 or lower ([#416](#)). If you are not already up-to-date, you should install the latest runtime version from <https://dotnet.microsoft.com/download/dotnet-core/3.1>

### 3.4.2 July 2020

Microsoft has published several newer Security Advisories regarding ASP.NET Core:

- Microsoft.ApplicationInsights.AspNetcore 2.12 was vulnerable to [CVE-2020-2224](#). We upgraded it to 2.14.
- Microsoft Security Advisory CVE-2020-0602 : ASP.NET Core Denial of Service Vulnerability, [#402](#) and
- Microsoft Security Advisory CVE-2020-0603 : ASP.NET Core Remote Code Execution Vulnerability, [#403](#). These affect applications running SignalR. Vonk does not use SignalR. Nevertheless we recommend to follow Microsoft's advice: For machines running .NET Core 3.1, you should download and install Runtime 3.1.1 or SDK 3.1.101 from <https://dotnet.microsoft.com/download/dotnet-core/3.1>
- Microsoft Security Advisory | MessagePack Denial of Service, [#405](#). This only affect applications using MessagePack, which Vonk does not use.

### 3.4.3 September 2019

Updates regarding previous Security Advisories:

- Please upgrade the ASP.NET Core runtime to at least version 2.2.7, from the [runtimes download page](#). This solves:
  - [#334](#)
  - [#373](#)
  - [#384](#)
  - [#385](#)
- [#295](#): Vonk has been upgraded to ASP.NET Core 2.2, and is therefore no longer vulnerable to this issue. It is nevertheless advised to run a publicly exposed Vonk *behind a proxy* or on an Azure Web App.
- [#335](#): no longer relevant to Vonk since it runs on ASP.NET Core 2.2

Microsoft has published several newer Security Advisories regarding ASP.NET Core:

- [#325](#): This is not applicable yet to Vonk, since it affects AspNetCoreModuleV2 and Vonk still works on AspNetCoreModule (implicitly V1). We will upgrade to V2 shortly though, so we advise you to install the latest AspNetCoreModulev2 anyway.
- [#359](#): Not relevant to Vonk, it does not use SignalR.

### 3.4.4 January 2019

Microsoft has published two Security Advisories regarding ASP.NET Core:

- If you run Vonk behind Internet Information Server (IIS), you may be vulnerable to “Microsoft Security Advisory CVE-2019-0548: ASP.NET Core Denial Of Service Vulnerability”. Refer to the related [Github issue #335](#) for details and the fix.
- When using older versions of some of the ASP.NET Core packages you may be vulnerable to “Microsoft Security Advisory CVE-2019-0564: ASP.NET Core Denial of Service Vulnerability”. Refer to the related [Github issue #334](#) for details. Vonk FHIR Server up until version 1.1.0 uses versions of the packages involved that are not affected (older than the vulnerable versions). In a future version we will upgrade beyond the vulnerable version upto secure versions. No action is required by the administrator of Vonk.

### 3.4.5 April 2018

Microsoft has published two Security Advisories regarding ASP.NET Core:

- If you run Vonk behind Internet Information Server (IIS), you may be affected by “Microsoft Security Advisory CVE-2018-0808: ASP.NET Core Denial Of Service Vulnerability”. Refer to the related [GitHub issue #294](#) for details and the fix.
- If you expose Vonk directly to the internet, or host it behind a proxy which does not validate or restrict host headers to known good values, you may be affected by “Microsoft Security Advisory CVE-2018-0787: ASP.NET Core Elevation Of Privilege Vulnerability”. Refer to the related [GitHub issue #295](#) for details and the correct way of hosting Vonk. This ‘host validating middleware’ mentioned by this issue is not a part of Vonk. We advise you to run a publicly exposed Vonk *behind a proxy* or on an Azure Web App.

## 3.5 How to upgrade Firely Server?

The process for upgrading Firely Server depends on whether you have a vanilla Firely Server, you added your own plugins or are running a Facade. This page describes the general process for each situation. Please refer to the [Release notes and upgrade management Firely Server](#) for details per each released version of Firely Server.

### 3.5.1 Upgrading Firely Server

#### Using the binary distribution

1. Download the latest version of Firely Server, see [Getting Started](#), and extract it to where you want it installed.
2. Copy your appsettings.instance.json and logsettings.instance.json files from the current installation to the new installation.
3. Check the [Release notes and upgrade management Firely Server](#) for any new settings that you may want to apply or change from their defaults.
4. Check the [Release notes and upgrade management Firely Server](#) for any actions that you need to take specifically for this upgrade.
5. Make sure the new installation can find the license file (see [License](#), general advice is to put the license file outside of the installation directory).
6. Create a backup of your current databases, both the main Resource database and the Administration database. See [Repository](#) to find the details on your configured database connection.

7. Stop the running instance of Firely Server (Ctrl + C if running from the console).
  8. Switch to the new installation directory and start Firely Server from there (`> dotnet ./Vonk.Server.dll`)
  9. Firely Server will now do several upgrade tasks, during which any web request will be responded to with 423 - Locked:
    1. If needed, an update is applied to the database structure.
    2. If Firely Server introduces a new version of the FHIR .NET API, Firely Server will load a new set of Conformance Resources from the `specification.zip` into the Administration database, for both FHIR STU3 and FHIR R4. In a specific case you can [prevent this step from happening](#).
  10. When Firely Server is done with the tasks above, it is again available to process requests.
  11. Check the log for warnings stating that you use obsolete settings. If so, adjust them and restart Firely Server.
- If anything went wrong, go back:
1. Stop the (new) running instance of Firely Server.
  2. Restore both databases from your backup.
  3. Switch to the old installation directory and start the old version of Firely Server from there (`> dotnet .\Vonk.Server.dll`)
  4. It should start as it did before you began the upgrade.
  5. Report the problem to the Firely Server helpdesk, see [Contact us](#).

You may be able to avoid the import of `specification.zip` if:

- The Administration database is in SQLite and
- You have not made alterations to the Administration API through the Web API.

In this case you can simply replace the old database (usually with the filename `vonkadmin.db`) with the one from the new installation directory (in `./data/vonkadmin.db`). Do so *before* you start the new Firely Server installation. Anything specified in [AdministrationImportOptions](#) will be re-imported into the new database.

## Using Docker

Revisit [Using Firely Server on Docker](#).

1. Stop the running container for Firely Server: `> docker stop vonk.server`.
2. Pull the latest image for Firely Server: `> docker pull firely/server` (versions 4.6 and older: `> docker pull simplifier/vonk`)
3. Check the [Release notes and upgrade management Firely Server](#) for any new settings that you may want to apply or change from their defaults, and apply that to the `environment` setting in the `docker-compose` file.
4. Check the [Release notes and upgrade management Firely Server](#) for any action that you need to take specifically for this upgrade.
5. Create a backup of your current databases, both the main Resource database and the Administration database. See [Repository](#) and your `docker-compose` file to find the details on where your databases are.
6. Start the new version (see [Using Firely Server on Docker](#) for the various commands to run the Firely Server container).
7. Firely Server will now do several upgrade tasks, during which any web request will be responded to with 423 - Locked:
  1. If needed, an update is applied to the database structure.

2. If Firely Server introduces a new version of the FHIR .NET API, Firely Server will load a new set of Conformance Resources from the specification.zip into the Administration database, for both FHIR STU3 and FHIR R4. In a specific case you can *prevent this step from happening*.
8. When Firely Server is done with the tasks above, it is again available to process requests.
9. Check the log for warnings stating that you use obsolete settings. If so, adjust them and restart Firely Server.

If anything went wrong, go back:

1. Stop the (new) running container of Firely Server.
2. Restore both databases from your backup.
3. Specify your previous image of Firely Server in the docker command or in the docker-compose file: `firely/server:<previous-version-tag>`
4. Start the container based on this previous image.
5. It should start as it did before you began the upgrade.
6. Report the problem to the Firely Server helpdesk, see *Contact us*.

### 3.5.2 Upgrading Plugins

Since a Plugin runs in the context of a Firely Server we advise you to start by upgrading your Firely Server, without loading your Plugin. Check the section on *Configuring the Firely Server Pipeline* to see how you can exclude your plugin from the pipeline.

**Attention:** We do not guarantee that a plugin built against version x.y.z of Firely Server can be run within a newer or older version as-is. Between minor versions recompilation is usually enough to update your plugin. Between major versions you should prepare for breaking changes in the public programming API. Sometimes we choose to apply such changes even on a minor version update, if we are fairly sure that you will not or only slightly be affected by it.

Upgrade the references in your plugin:

1. Open the source code of your plugin, and open the project file (yourplugin.csproj).
2. Change the references to the Firely Server.\* packages to the version that you want to upgrade to.
3. Build and check the errors.
4. Check the list of breaking changes for the new Firely Server version in the *Release notes and upgrade management Firely Server*. Applying the changes should fix the errors.
5. With some releases Firely Server is also upgraded to a newer version of the Firely .NET SDK. That is then mentioned in the release notes. If this is the case, also check the *SDK release notes* for breaking changes.
6. Still errors? Maybe we have overlooked a change. Please report it to us, see *Contact us*. And if it is easy to fix - do so :-)
7. Build and publish your plugin.
8. Put the resulting dll's in the plugin directory of the new installation of Firely Server.
9. Re-include your plugin in the pipeline.
10. (Re)start Firely Server and test the working of your plugin.

### 3.5.3 Upgrading Facades

A Facade implementation is technically also a plugin, but one that only adds repository access services. For this it makes no sense to try to run Firely Server without the Facade as is described for plugins. So start with upgrading the references right away.

Especially for Facades to relational databases: match the version of EntityFrameworkCore that the new version of Firely Server is using. Check the list of changes to see whether we upgraded.

### 3.5.4 Migrations

Sometimes, an upgrade of the database schema is needed in order to introduce performance improvements or enable new features. Depending on the database that is used, Firely Server provides different techniques to migrate the data. Most migrations can be performed automatically whereas in some cases, manual migrations are preferable.

One of these cases is a scenario in which Firely Server operates at scale, handling billions of resources. In such a scenario, manual migrations are preferable as they might take longer than expected and come with challenges. In this section, we will elaborate on some of these challenges and how to mitigate them.

#### MongoDB

For Firely Server instances running on MongoDB, we consider databases with more than 500GB as large. Migrations for these databases take a long time which is why we introduced external migration scripts with Firely Server 4.5.0. These external scripts are written in JavaScript and allow to exert more control at the cost of some convenience.

**The migration scripts usually perform two operations in this order:**

1. Migrate the existing data using the `updateMany()` operator.
2. Update the system information document to the next version.

While migrating large databases manually via SSH it is always a risk that the connection with the remote computer breaks and the SSH session terminates. Normally, this would also mean that all programs started within that SSH session will also be terminated including the migration. To prevent this from happening, we recommend using the `screen` tool to perform the migration. Using it allows programs to continue running even if the SSH connection breaks.

1. Install the `screen` tool or check whether it is installed with `screen --version`
2. Start a screen session: `screen`
3. Execute the migration script in this screen session
4. In case the SSH connection breaks, you can connect via SSH again and get access to the running migration process using the command `screen -r`
5. Wait until the migration is done and restart Firely Server

In case the migration actually times out or is interrupted, the `updateMany()` operator will continue running in the background. The update of the schema version, however, will not succeed and needs to be performed manually. Follow these steps to do so:

1. Connect to your MongoDB host.
2. Open a MongoDB shell by executing `mongo` from the command line.
3. Switch to the `vonk` database: use `vonkdata`

4. Check whether the current migration is still running with: `db.currentOp()`. You will likely see multiple operations in this command's output. Look for an operation that resembles the commands included in the `try {} catch {}` block in the migration's JavaScript file. If any such operation is still listed, the migration is still running.
5. Wait until the operation associated with the migration has finished.
6. Afterwards, retrieve the system document id by executing:

```
db.vonkentries.findOne(  
  {  
    "LatestMigration": { $exists: true }  
  }  
)
```

7. Note down the value of the `_id` field, for example `ObjectId("61bc7dab260c691f4c0f78d4")`
8. Open the JavaScript migration script. Note down the value of the next FS and migration version. You will find these values as constants at the top of the migration script or in a `coll.updateOne()` command at the end of the script.
9. Open a mongo shell on your host and exchange `ObjectId`, `VonkVersion` and `LatestMigration` to the previously noted values:

```
db.vonkentries.updateOne(  
  {  
    _id: ObjectId("61bc7dab260c691f4c0f78d4")  
  },  
  {  
    "$set": {  
      VonkVersion: "4.9.0",  
      LatestMigration: 22  
    }  
  }  
)
```

10. Execute this command in your mongo shell

That's it! The migration of the data is finished and the document containing the system information is updated accordingly. You should now be able to start the new version of Firely Server.

## 3.6 Firely Server Roadmap

This page lists the features and changes we have planned for the foreseeable future. This planning is volatile, so changes will happen. We will update this page accordingly. You are also very welcome to provide input to us on features or fixes that you would like to see prioritized.

Disclaimer: No rights can be derived from this roadmap.

### 3.6.1 2022

#### Q2

- Addition of various performance improvements and features to our MongoDB implementation which are currently only available for SQL Server
  - Support for transactions

#### Q3

- Check compliance with CMS ImplementationGuides
- Bulk data export as a separate service - deploy and scale it independently of Firely Server
- New Firely Validator as a separate service - deploy and scale it independently of Firely Server
- Provide production-ready IdentityServer instance bundled with Firely Server

#### Q4

- Enhanced support for subscriptions
- R4B Support





## FREQUENTLY ASKED QUESTIONS

... and known issues.

### 4.1 Conflicting resources upon import

When importing specification.zip for R4, Firely Server will report errors like these:

```
Artifact C:\Users\<user>\AppData\Local\Temp\FhirArtifactCache-1.5.0-Hl7.Fhir.R4.
↪Specification\specification_Fhir4_0\dataelements.xml could not be imported. Error_
↪message: Found multiple conflicting resources with the same resource uri identifier.

Url: http://hl7.org/fhir/StructureDefinition/de-Quantity.value
  File: C:\Users\Christiaan\AppData\Local\Temp\FhirArtifactCache-1.5.0-Hl7.Fhir.R4.
↪Specification\specification_Fhir4_0\dataelements.xml
  File: C:\Users\Christiaan\AppData\Local\Temp\FhirArtifactCache-1.5.0-Hl7.Fhir.R4.
↪Specification\specification_Fhir4_0\dataelements.xml
  File: C:\Users\Christiaan\AppData\Local\Temp\FhirArtifactCache-1.5.0-Hl7.Fhir.R4.
↪Specification\specification_Fhir4_0\dataelements.xml
```

The error message is actually correct, since there *are* duplicate fullUrls in dataelements.xml in the specification. This has been reported in [Jira issue FHIR-25430](#).

### 4.2 Searchparameter errors for composite parameters

When importing specification.zip for various FHIR versions, Firely Server will report errors like these:

```
Composite SearchParameter 'CodeSystem.context-type-quantity' doesn't have components.
```

A searchparameter of type 'composite' should define which components it consists of. Firely Server checks whether all the components of such a composite searchparameter are present. If no components are defined at all - that is, SearchParameter.component is empty - it will display this error. This indicates an error in the definition of the searchparameter and should be solved by the author of it.

However, the implementation of this check seems to have an error so too many composite parameters are reported as faulty. We will address this issue in the next release.

## 4.3 .NET SDK not found

Since version 4.0 Vonk was renamed to Firely Server, including the main entrypoint. It changed from `vonk.server.dll` to `firely.server.dll`.

If you now still run `dotnet vonk.server.dll` on .NET runtime 3.1 it will state this error:

```
It was not possible to find any installed .NET Core SDKs
Did you mean to run .NET Core SDK commands? Install a .NET Core SDK from: https://aka.ms/
dotnet-download
```

This is very misleading. The actual error is that you probably tried to run `dotnet vonk.server.dll` but this dll no longer exists.

The same error can happen if you have built a Docker image of your own with `dotnet vonk.server.dll` as entry-point.

.NET 5 fixed this and more clearly states that the dll is missing.

## 4.4 Homepage takes long to load

The html homepage that is provided with Firely Server may take a long time to load, even though the server seems fully up and running.

The homepage consumes the bootstrap.js library from a CDN. The delay may be caused by a firewall slowing down that download.

The remedy is to disable the homepage in the pipelinesettings:

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        "Vonk.Core",
        ...,
        // "Vonk.UI.Demo", <-- disable this one
        ...
      ],
    },
    {
      "Path": "/administration",
      ...
    }
  ]
}
```

## CONFIGURING FIRELY SERVER

In this section we assume you have downloaded and installed the Firely Server binaries, and have obtained a license file. If not, please see the [Getting Started](#) and follow the steps there first.

The steps you followed to get started will provide you with a basic Firely Server, that runs on a standard port and keeps the data in a SQLite database.

If you need to adjust the port, or want to use a MongoDB or SQL database you can configure Firely Server by adjusting the [Firely Server settings](#).

If you want to change the way Firely Server logs its information, you can adjust the [Log settings](#).

### 5.1 Firely Server settings

Firely Server settings are controlled in json configuration files called `appsettings(*).json`. The possible settings in these files are all the same and described below. The different files are read in a hierarchy so you can control settings on different levels. All `appsettings` files are in the Firely Server distribution directory, next to `Firely.Server.dll`. We go through all the sections of this file and refer you to detailed pages on each of them.

You can also control [Firely Server settings with Environment Variables](#).

Changes to the settings require a restart of Firely Server.

#### 5.1.1 Hierarchy of settings

Firely Server reads its settings from these sources, in this order:

**appsettings.default.json**

Installed with Firely Server, contains default settings and a template setting if no sensible default is available.

**appsettings.json**

You can create this one for your own settings. Because it is not part of the Firely Server distribution, it will not be overwritten by a next Firely Server version.

**environment variables**

See [Firely Server settings with Environment Variables](#).

**appsettings.instance.json**

You can create this one to override settings for a specific instance of Firely Server. It is not part of the Firely Server distribution. This file is especially useful if you run multiple instances on the same machine.

Settings lower in the list override the settings higher in the list (think CSS, if you're familiar with that).

**Warning:** JSON settings files can have arrays in them. The configuration system can NOT merge arrays. So if you override an array value, you need to provide all the values that you want in the array. In the Firely Server settings this is relevant for e.g. `Validation.AllowedProfiles` and for the `PipelineOptions`.

**Note:** By default in ASP.NET Core, if on a lower level the array has more items, you will still inherit those extra items. We fixed this in Firely Server, an array will always overwrite the complete base array. To nullify an array, add the value with an array with just an empty string in it:

```
"PipelineOptions": {
  "Branches": [
    {
      "Path": "myroot",
      "Exclude": [""]
    }
  ]
}
```

This also means you cannot override a single array element with an environment variable. (Which was tricky anyway - relying on the exact number and order of items in the original array.)

### 5.1.2 Changing the settings

In general you do not change the settings in `appsettings.default.json` but create your own overrides in `appsettings.json` or `appsettings.instance.json`. That way your settings are not overwritten by a new version of Firely Server (with a new `appsettings.default.json` therein), and you automatically get sensible defaults for any new settings introduced in `appsettings.default.json`.

#### Settings after first install

After you installed Firely Server (see [Getting Started](#)), either:

- copy the `appsettings.default.json` to `appsettings[.instance].json` and remove settings that you do not intend to alter, or
- create an empty `appsettings[.instance].json` and copy individual parts from the `appsettings.default.json` if you wish to adjust them.

Adjust the new `appsettings[.instance].json` to your liking using the explanation below.

When running [Firely Server on Docker](#) you probably want to adjust the settings using the [Environment Variables](#).

## Settings after update

If you install the binaries of an updated version of Firely Server, you can:

- copy the new binaries over the old ones, or
- deploy the new version to a new directory and copy the `appsettings[.instance].json` over from the old version.

In both cases, check the *Release notes and upgrade management Firely Server* to see if settings have changed, or new settings have been introduced. If you want to adjust a changed / new setting, copy the relevant section from `appsettings.default.json` to your own `appsettings[.instance].json` and then adjust it.

## Commenting out sections

JSON formally has no notion of comments. But the configuration system of ASP.NET Core (and hence Firely Server) accepts double slashes just fine:

```
"Administration": {
  "Repository": "SQLite", //Memory / SQL / MongoDB
  "SqlDbOptions": {
    "ConnectionString": "connectionstring to your Firely Server Admin SQL Server,
    ↳ database (SQL2012 or newer); Set MultipleActiveResultSets=True",
    "SchemaName": "vonkadmin",
    "AutoUpdateDatabase": true,
    "MigrationTimeout": 1800 // in seconds
    // "AutoUpdateConnectionString" : "set this to the same database as
    ↳ 'ConnectionString' but with credentials that can alter the database. If not set,
    ↳ defaults to the value of 'ConnectionString'"
  },
}
```

This will ignore the `AutoUpdateConnectionString`.

### 5.1.3 Log of your configuration

Because the hierarchy of settings can be overwhelming, Firely Server logs the resulting configuration. To enable that, the loglevel for `Vonk.Server` must be `Information` or more detailed. That is set for you by default in `logsettings.default.json`. Refer to *Log settings* for information on setting log levels.

### 5.1.4 Administration

```
"Administration": {
  "Repository": "SQLite", //Memory / SQL / MongoDB are other options, but SQLite is
  ↳ advised.
  "MongoDbOptions": {
    "ConnectionString": "mongodb://localhost/vonkadmin",
    "EntryCollection": "vonkentries"
  },
  "SqlDbOptions": {
    "ConnectionString": "connectionstring to your Firely Server Admin SQL Server,
    ↳ database (SQL2012 or newer); Set MultipleActiveResultSets=True",
    "SchemaName": "vonkadmin",
  },
}
```

(continues on next page)

(continued from previous page)

```

    "AutoUpdateDatabase": true,
    "MigrationTimeout": 1800 // in seconds
    // "AutoUpdateConnectionString" : "set this to the same database as
    ↪ 'ConnectionString' but with credentials that can alter the database. If not set, ↪
    ↪ defaults to the value of 'ConnectionString'"
  },
  "SQLiteDbOptions": {
    "ConnectionString": "Data Source=./data/vonkadmin.db",
    "AutoUpdateDatabase": true
  },
  "Security": {
    "AllowedNetworks": [ "::1" ], // i.e.: ["127.0.0.1", "::1" (ipv6 localhost), "10.1.
    ↪ 50.0/24", "10.5.3.0/24", "31.161.91.98"]
    "OperationsToBeSecured": [ "reindex", "reset", "preload" ]
  }
},

```

The Administration section is to *Configure the Administration API* and its repository.

### 5.1.5 License

```

"License": {
  "LicenseFile": "firelyserver-trial-license.json",
  "RequestInfoFile": "./.vonk-request-info.json",
  "WriteRequestInfoFileInterval": 15 // in minutes
}

```

The *Getting Started* explains how to obtain a licensefile for Firely Server. Once you have it, put the path to it in the LicenseFile setting. Note that in json you either use forward slashes (/) or double backward slashes (\\) as path separators.

---

**Note:** It is also possible to supply a license via an environment variable. This functionality is handy when Firely Server is running within a Docker container. See *Providing license via an environment variable* for details.

---

Other settings:

- RequestInfoFile sets the location of the file with request information. This file will be used in future releases.
- WriteRequestInfoFileInterval sets the time interval (in minutes) to write aggregate information about processed requests to the RequestInfoFile.

### 5.1.6 Repository

```
"Repository": "SQLite", //Memory / SQL / MongoDB
```

1. Repository: Choose which type of repository you want. Valid values are:

1. Memory
2. SQL, for Microsoft SQL Server
3. SQLite
4. MongoDB

#### Memory

```
"MemoryOptions": {
  "SimulateTransactions": "false"
},
```

Refer to *Using the In-Memory storage* for configuring the In-Memory storage.

#### MongoDB

```
"MongoDbOptions": {
  "ConnectionString": "mongodb://localhost/vonkdata",
  "EntryCollection": "vonkentries"
},
```

Refer to *Using MongoDB* for configuring the connection to your MongoDB databases.

#### SQL

```
"SqlDbOptions": {
  "ConnectionString": "connectionstring to your Firely Server SQL Server database.
↳(SQL2012 or newer); Set MultipleActiveResultSets=True",
  "SchemaName": "vonk",
  "AutoUpdateDatabase": true,
  "MigrationTimeout": 1800 // in seconds
  //"AutoUpdateConnectionString" : "set this to the same database as 'ConnectionString
↳' but with credentials that can alter the database. If not set, defaults to the value.
↳of 'ConnectionString'"
},
```

Refer to *Using SQL server* for configuring access to your SQL Server databases.

### SQLite

```
"SQLiteDbOptions": {  
  "ConnectionString": "Data Source=./data/vonkdata.db",  
  "AutoUpdateDatabase": true  
},
```

Refer to *Using SQLite* for configuring access to your SQLite Server databases.

### 5.1.7 http and https

```
"Hosting": {  
  "HttpPort": 4080,  
  //"HttpsPort": 4081, // Enable this to use https  
  //"CertificateFile": "<your-certificate-file>.pfx", //Relevant when HttpsPort is  
  ➔present  
  //"CertificatePassword" : "<cert-pass>" // Relevant when HttpsPort is present  
  //"PathBase": "<subpath-to-firely-server>",  
  "ClientCertificateMode": "NoCertificate" // NoCertificate, AllowCertificate,  
  ➔RequireCertificate  
},
```

Refer to *Configure http and https* for enabling https and adjusting port numbers. The *PathBase* enables the option to specify a path as part of root path (See *PathBase middleware* for more information). The *ClientCertificateMode* will instruct Firely Server to request or require a TLS client certificate (See *ASP .NET Core - Client Certificates* for more information).

### 5.1.8 Validation

```
"Validation": {  
  "Parsing": "Permissive", // Permissive / Strict  
  "Level": "Off", // Off / Core / Full  
  "AllowedProfiles": []  
},
```

Refer to *Validating incoming resources*.

### 5.1.9 Search size

```
"BundleOptions": {  
  "DefaultCount": 10,  
  "MaxCount": 50,  
  "DefaultSort": "-_lastUpdated"  
},
```

The Search interactions returns a bundle with results. Users can specify the number of results that they want to receive in one response with the `_count` parameter.

- `DefaultCount` sets the number of results if the user has not specified a `_count` parameter.



- **MaxCount** sets the number of results in case the user specifies a `_count` value higher than this maximum. This is to protect Firely Server from being overloaded.
- **DefaultCount** should be less than or equal to **MaxCount**
- **DefaultSort** is what search results are sorted on if no sort order is specified in the request. If a sort order is specified, this is still added as the last sort clause.

### 5.1.10 Batch and transaction

```
"BatchOptions": {
  "MaxNumberOfEntries": 100
},
```

This will limit the number of entries that are accepted in a single Batch or Transaction bundle.

**Note:** This setting has been moved to the **SizeLimits** setting as of Firely Server (Vonk) version 0.7.1, and the logs will show a warning that it is deprecated when you still have it in your appsettings file.

### 5.1.11 Response options

```
"HttpOptions": {
  "DefaultResponseType": "application/fhir+json"
}
```

- If no mediatype is specified in an **Accept** header, use the **DefaultResponseType**.
- Options are `application/fhir+json` or `application/fhir+xml`
- Firely Server will attach the mimetype parameter `fhirVersion` based on the FHIR version that is requested (see *Multiple versions of FHIR*).

### 5.1.12 Protect against large input

```
"SizeLimits": {
  "MaxResourceSize": "1MiB",
  "MaxBatchSize": "5MiB",
  "MaxBatchEntries": 150
},
```

- **MaxResourceSize** sets the maximum size of a resource that is sent in a create or update.
- **MaxBatchSize** sets the maximum size of a batch or transaction bundle. (Note that a POST `http(s)://<firely-server-endpoint>/Bundle` will be limited by **MaxResourceSize**, since the bundle must be processed as a whole then.)
- **MaxBatchEntries** limits the number of entries that is allowed in a batch or transaction bundle.
- The values for **MaxResourceSize** and **MaxBatchSize** can be expressed in b (bytes, the default), kB (kilobytes), KiB (kibibytes), MB (megabytes), or MiB (mebibytes). Do not put a space between the amount and the unit.

### 5.1.13 SearchParameters and other Conformance Resources

```
"AdministrationImportOptions": {
  "ImportDirectory": "./vonk-import",
  "ImportedDirectory": "./vonk-imported", //Do not place ImportedDirectory *under*
↳ ImportDirectory, since an import will recursively read all subdirectories.
  "SimplifierProjects": [
    {
      "Uri": "https://stu3.simplifier.net/<your-project>",
      "UserName": "Simplifier user name",
      "Password": "Password for the above user name",
      "BatchSize": 20
    }
  ]
}
```

See *Managing Conformance Resources* and *Custom Search Parameters*.

### 5.1.14 Cache of Conformance Resources

```
"Cache": {
  "MaxConformanceResources": 5000
}
```

Firely Server caches StructureDefinitions and other conformance resources that are needed for (de)serialization and validation in memory. If more than MaxConformanceResources get cached, the ones that have not been used for the longest time are discarded. If you frequently encounter a delay when requesting less used resource types, a larger value may help. If you are very restricted on memory, you can lower the value.

### 5.1.15 Reindexing for changes in SearchParameters

```
"ReindexOptions": {
  "BatchSize": 100,
  "MaxDegreeOfParallelism": 10
},
```

See *Re-index Configuration*.

### 5.1.16 Restrict supported resources and SearchParameters

```
"SupportedModel": {
  "RestrictToResources": [ "Patient", "Observation" ],
  "RestrictToSearchParameters": ["Patient.active", "Observation.patient", "Resource._id",
↳ "StructureDefinition.url"],
  "RestrictToCompartments": ["Patient"]
},
```

By default, Firely Server supports all ResourceTypes, SearchParameters and CompartmentDefinitions from the specification. They are loaded from the *specification.zip*. If you want to limit support, you can do so with the configuration

above. This is primarily targeted towards Facade builders, because they have to provide an implementation for everything that is supported.

Be aware that:

- support for `_type` and `_id` must not be disabled
- the Administration API requires support for the 'url' SearchParameter on the conformance resourcetypes
- this uses the search parameter names, not the path within the resource - so for example to specify `Patient.address.postalCode` as a supported location, you'd use `"Patient.address-postalcode"`.

### 5.1.17 Enable or disable interactions

By default, the value `SupportedInteractions` contains all the interactions that are implemented in Firely Server. But you can disable interactions by removing them from these lists.

```
"SupportedInteractions": {
  "InstanceLevelInteractions": "read, vread, update, delete, history, conditional_
↪delete, conditional_update, $validate",
  "TypeLevelInteractions": "create, search, history, $validate, $snapshot, conditional_
↪create",
  "WholeSystemInteractions": "capabilities, batch, transaction, history, search,
↪$validate"
},
```

If you implement a custom operation in a plugin, you should also add the name of that operation at the correct level. E.g. add `$convert` to `TypeLevelInteractions` to allow `<base>/<resourcetype>/$convert`.

### 5.1.18 Subscriptions

```
"SubscriptionEvaluatorOptions": {
  "Enabled": true,
  "RepeatPeriod": 20000,
  "SubscriptionBatchSize" : 1
},
```

See *Subscriptions*.

### 5.1.19 Information model

Firely Server supports the use of multiple information models (currently FHIR STU3 and R4) simultaneously. The `InformationModel` section contains the related settings. By default, Firely Server serves both versions from the root of your web service, defaulting to R4 when the client does not use `Accept` or `_format` to specify either one. Mapping a path or a subdomain to a specific version creates an additional URI serving only that particular version.

```
"InformationModel": {
  "Default": "Fhir4.0", // For STU3: "Fhir3.0". Information model to use when none is_
↪specified in either mapping, the _format parameter or the ACCEPT header.
  "Mapping": {
    "Mode": "Off"
    //"Mode": "Path", // yourserver.org/r3 => FHIR STU3; yourserver.org/r4 => FHIR R4
    //"Map": {
```

(continues on next page)

(continued from previous page)

```

    // "/R3": "Fhir3.0",
    // "/R4": "Fhir4.0"
  //}
  // "Mode": "Subdomain", // r3.yourserver.org => FHIR STU3; r4.yourserver.org => FHIR_
↪R4
  // "Map":
  // {
  //   "r3": "Fhir3.0",
  //   "r4": "Fhir4.0"
  // }
}
},

```

See *Multiple versions of FHIR*.

## 5.1.20 Patient Everything Operation

```

"PatientEverythingOperation": {
  "AdditionalResources": [ "Organization", "Location", "Substance", "Medication", "Device
↪" ] // included referenced resources, additional to the Patient compartment resources
},

```

The Patient \$everything operation returns all resources linked to a patient that are listed in the Compartment Patient. This section allows you to define additional resources that will be included in the resulting searchset bundle.

See *Patient \$everything*.

## 5.1.21 FHIR Capabilities

```

"FhirCapabilities": {
  "ConditionalDeleteOptions": {
    "ConditionalDeleteType": "Single", // Single or Multiple,
    "ConditionalDeleteMaxItems": 1
  }
},

```

See *Create, read, update, patch, delete*.

## 5.1.22 History size

```

"HistoryOptions": {
  "MaxReturnedResults": 100
}

```

See *History*.

The settings for authorization with SMART on FHIR are covered in *Configuration*.

### 5.1.23 Configuring the Firely Server Pipeline

You can add your own plugins to the Firely Server pipeline, or control which of the standard Firely Server plugins are used for your Firely Server, by changing the PipelineOptions.

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",
        "Vonk.Fhir.R4",
        // etc.
      ],
      "Exclude": [
      ]
    },
    {
      "Path": "/administration",
      "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",
        "Vonk.Fhir.R4",
        // etc.
      ],
      "Exclude": [
        "Vonk.Core.Operations"
      ]
    }
  ]
}
```

It is possible to disable a specific information model by removing Vonk.Fhir.R3 or Vonk.Fhir.R4 from the pipeline

Please note the warning on merging arrays in *Hierarchy of settings*.

See *Firely Server Plugins* for more information and an example custom plugin.

## 5.2 Firely Server settings with Environment Variables

### 5.2.1 Environment Variables for appsettings

All the settings in *Firely Server settings* can be overridden by environment variables on your OS. This can be useful if you want to deploy Firely Server to several machines, each having their own settings for certain options. For *Using Firely Server on Docker* using environment variables in the docker-compose file is currently the only way to pass settings to the container. Or if you don't want a database password in the appsettings.json file.

The format for the environment variables is:

```
VONK_<setting_level_1>[:<setting_level_n>]*
```

So you start the variable name with the prefix 'VONK\_', and then follow the properties in the json settings, separating each level with a colon ':'. Some examples:

appsettings.json:

```
"Repository" : "SQL"
```

environment variable:

```
VONK_Repository=SQL
```

To access an embedded value, using the '.' separator:

appsettings.json:

```
"Administration" : {  
  "SqlDbOptions" : {  
    "ConnectionString" : "<some connectionstring>"  
  }  
}
```

environment variable:

```
VONK_Repository:SqlDbOptions:ConnectionString=<some connectionstring>
```

To access an array item, use 0-based indexing:

```
VONK_PipelineOptions:Branches:0:Exclude:0=Vonk.Repository.Memory  
VONK_PipelineOptions:Branches:0:Exclude:1=Vonk.Repository.Sql
```

### Arrays in Environment Variables

Sometimes the appsettings allow for an array of values, e.g. in the setting for `AllowedProfiles` in *Validating incoming resources*. You can address them by appending an extra colon and an index number.

appsettings.json:

```
"Validation": {  
  "ValidateIncomingResources": "true",  
  "AllowedProfiles":  
  [  
    "http://hl7.org/fhir/StructureDefinition/daf-patient",  
    "http://hl7.org/fhir/StructureDefinition/daf-allergyintolerance"  
  ]  
}
```

environment variables:

```
VONK_Validation:ValidateIncomingResources=true VONK_Validation:AllowedProfiles:0=http://hl7.org/fhir/StructureDefinition  
patient VONK_Validation:AllowedProfiles:1=http://hl7.org/fhir/StructureDefinition/daf-  
allergyintolerance
```

## 5.2.2 Log settings with Environment Variables

You can control the *Log settings* with Environment Variables the same way as the *Environment Variables for appsettings* above. The difference is in the prefix. For the log settings we use 'VONKLOG\_'.

logsettings.json

```
"Serilog": {
  "MinimumLevel": {
    "Override": {
      "Vonk.Configuration": "Information",
```

environment variable:

VONKLOG\_Serilog:MinimumLevel:Override:Vonk.Configuration=Information

## 5.2.3 Audit log settings with Environment Variables

You can control the *Audit log file configuration* with Environment Variables the same way as the *Environment Variables for appsettings* above. The difference is in the prefix. For the log settings we use 'VONKAUDITLOG\_'.

audit.logsettings.json

```
"AuditLog": {
  "WriteTo": [
    {
      "Name": "File",
      "Args": {
        "path": "./audit/AuditLog.log"
```

environment variable:

VONKAUDITLOG\_AuditLog:WriteTo:0:Args:path=./other/directory/AuditLog.log

## 5.2.4 Changing Environment Variables on Windows

In Windows you can change the Environment Variables with Powershell or through the UI. Based on the first example above:

- In Powershell run: > \$env:VONK\_Repository="SQL"
- or go to your *System*, open the *Advanced system settings* → *Environment variables* and create a new variable with the name VONK\_Repository and set the value to "SQL" (you don't need to enter the quotes here).

## 5.3 Configure the Administration API

This configuration is part of *Firely Server settings*.

```
"Administration": {
  "Repository": "SQLite", //Memory / SQL / MongoDB
  "MongoDbOptions": {
    "ConnectionString": "mongodb://localhost/vonkadmin",
```

(continues on next page)

(continued from previous page)

```

    "EntryCollection": "vonkentries"
  },
  "SqlDbOptions": {
    "ConnectionString": "connectionstring to your Firely Server Admin SQL Server_
↪ database (SQL2012 or newer); Set MultipleActiveResultSets=True",
    "SchemaName": "vonkadmin",
    "AutoUpdateDatabase": true,
    "MigrationTimeout": 1800 // in seconds
    // "AutoUpdateConnectionString" : "set this to the same database as 'ConnectionString
↪ ' but with credentials that can alter the database. If not set, defaults to the value_
↪ of 'ConnectionString'"
  },
  "SQLiteDbOptions": {
    "ConnectionString": "Data Source=./data/vonkadmin.db",
    "AutoUpdateDatabase": true,
    "MigrationTimeout": 1800 // in seconds
  },
  "Security": {
    "AllowedNetworks": [ "::1" ], // i.e.: ["127.0.0.1", "::1" (ipv6 localhost), "10.1.
↪ 50.0/24", "10.5.3.0/24", "31.161.91.98"]
    "OperationsToBeSecured": [ "reindex", "reset", "preload" ]
  }
},

```

### 5.3.1 Choosing your storage

1. Repository: Choose which type of repository you want. Valid values are:
  1. Memory
  2. SQL
  3. SQLite
  4. MongoDB
1. MongoDBOptions: Use these with "Repository": "MongoDb", see *Using MongoDB* for details.
2. SqlDbOptions: Use these with "Repository": "SQL", see *Using SQL server* for details.
3. SQLiteDbOptions: Use these with "Repository": "SQLite", see *Using SQLite* for details.

### 5.3.2 Limited access

1. Security: You can restrict access to the operations listed in OperationsToBeSecured to only be invoked from the IP addresses listed in AllowedNetworks.
  - Operations that can be secured are:
    - reindex (see *Re-indexing for new or changed SearchParameters*)
    - reset (see *Reset the database*)
    - preload (see *Preloading resources*)
    - StructureDefinition (restrict both read and write)



- SearchParameter (restrict both read and write)
  - ValueSet (restrict both read and write)
  - CodeSystem (restrict both read and write)
  - CompartmentDefinition (restrict both read and write)
  - Subscription: (restrict both read and write, see *Subscriptions*)
- The AllowedNetworks have to be valid IP addresses, either IPv4 or IPv6, and masks are allowed.

## 5.4 Managing Conformance Resources

Firely Server uses *Conformance Resources* along with some *Terminology Resources* for various operations:

- SearchParameter: For indexing resources and evaluating *search* interactions.
- StructureDefinition: For *snapshot generation*, and of course – along with ValueSet and CodeSystem – for *validation*.
- CompartmentDefinition: For *access control* and Compartment *Search*.
- ValueSet and CodeSystem: For *Terminology services* operations.
- StructureMap and ConceptMap: for mapping

You can control the behaviour of Firely Server for these interactions by loading resources of these types into Firely Server. There are two ways of doing this:

1. With regular FHIR interactions (create, update, delete) on the *Firely Server Administration API*.
2. With the *Import of Conformance Resources*.

No matter which method you use, all Conformance resources are persisted in the Administration API database (see *Configure the Administration API* for configuring that database), and available through the Administration API endpoint (<firely-server-endpoint>/administration)

For each resourcetype the base profile is listed in the CapabilityStatement under CapabilityStatement.rest.resource.profile and (since FHIR R4) all the other profiles are listed under CapabilityStatement.rest.resource.supportedProfile. So by requesting the *CapabilityStatement* you can easily check whether your changes to the StructureDefinitions were correctly processed by Firely Server.

**Attention:** Please be aware that Conformance Resources have to have a **unique canonical url** within the FHIR Version they are loaded, in their url element. Firely Server does not allow you to POST two conformance resources with the same canonical url. For SearchParameter resources, the combination of base and code must be unique.

**Attention:** Creates or updates of **SearchParameter** resources should be followed by a *re-index*.

Before you delete a SearchParameter, be sure to remove it from the index first, see the *exclude* parameter in *re-index*.

Changes to the other types of resources have immediate effect.

**Attention:** A StructureDefinition can only be posted in the context of a FHIR Version that matches the StructureDefinition.fhirVersion. See *Multiple versions of FHIR*.

**Note:** The import process will retain any id that is already in the resource, or assign a new id if there is none in the resource. For FHIR versions other than STU3, a postfix is appended to the id to avoid collisions between FHIR versions. See also *Conformance resources*.

---

### 5.4.1 Import of Conformance Resources

The import process of conformance resources runs on every startup of Firely Server, and *on demand*.

The process uses these locations on disk:

- ImportDirectory;
- ImportedDirectory;
- a read history in the file `.vonk-import-history.json`, written in ImportedDirectory.

**Attention:** Please make sure that the Firely Server process has write permission on the ImportedDirectory.

The process follows these steps for each FHIR version (currently STU3 and R4, and experimentally for R5)

1. Load the *Default Conformance Resources*, if they have not been loaded before.
2. Load the *Errata to the specification*, if they have not been loaded before.
3. *Load Conformance Resources from disk*. After reading, the read files are registered in the read history.
4. *Load Conformance Resources from simplifier.net*. After reading, the project is registered in the read history. Subsequent reads will query only for resources that have changed since the last read.

Loading the conformance resources from the various sources can take some time, especially on first startup when the *Default Conformance Resources* have to be imported. During the import Firely Server will respond with 423 'Locked' to every request to avoid storing or retrieving inconsistent data.

The read history keeps a record of files that have been read, with an MD5 hash of each. If you wish to force a renewed import of a specific file, you should:

- manually edit the read history file and delete the entry about that file;
- provide the file again in the ImportDirectory (if you deleted it previously - Vonk does not delete it).

### 5.4.2 Retain the import history

If you run the Administration database on SQL Server or MongoDB it is important to *retain* the `.vonk-import-history` file. This means that if you run Firely Server on something stateless like a Kubernetes pod, or a webapp service, you need to attach file storage on which to store this file. If you do not do that, Firely Server will import all the conformance resources *on every start*.

### 5.4.3 Running imports with multiple instances

If you run multiple instances of Firely Server each will have its own `/administration` pipeline. So you need to make sure that only 1 instance will perform the import. The import at startup will happen when:

- we upgraded to a new version on the FHIR .NET API (always mentioned in the [releasenotes](#))
- you add new resources to the `ImportDirectory`
- resources retrieved from Simplifier are renewed.

To ensure that only one instance runs the import you can do two things:

1. Make sure only 1 instance is running:
  1. Stop Firely Server
  2. Scale down to 1 instance
  3. Upgrade Firely Server (by referring to a newer image, or installing newer binaries)
  4. Start Firely Server
  5. Let it do the import
  6. Then scale back up to multiple instances.
2. Exclude the namespace `Vonk.Administration.Api.Import` from the `PipelineOptions` in branch `administration` on all but one instance.

If you want to use the manual import (`<url>/administration/import`) you are advised to apply solution nr. 1 above. In the second solution the call may or may not end up on an instance having the Import functionality.

We are aware that this can be a bit cumbersome. On the [Firely Server Roadmap](#) is therefore the story to host the Administration API in its own microservice.

### 5.4.4 Default Conformance Resources

Firely Server comes with the `specification.zip` file from the HL7 FHIR API. It contains all the Conformance resources from the specification. These are loaded and used for validation and snapshot generation by default.

Some of the conformance resources (especially `SearchParameters`) contain errors in the core specification. We try to correct all errors in [Errata to the specification](#). You can also override them yourself by:

- updating them through the administration api, as described below;
- providing an altered version in the `ImportDirectory`, with the same id and canonical url.

**Attention:** The Core Specification provides almost 4000 Conformance Resources. Depending on the machine it may take a few minutes to load and index them.

### 5.4.5 Load Conformance Resources from disk

Firely Server can read SearchParameter and CompartmentDefinition resources from a directory on disk at startup. The AdministrationImportOptions in the *Firely Server settings* control from which directory resources are loaded:

```
"AdministrationImportOptions": {
  "ImportDirectory": "<path to the directory you want to import from, default ./vonk-
↳import>",
  "ImportedDirectory": "<path to the directory where imported files are moved to,↳
↳default ./vonk-imported>"
},
```

#### ImportDirectory

All files and zip files will be read, and any conformance resources in them will be imported. By default, STU3 is assumed. If you have R4 conformance resources, place them in a sibling directory that has the same name as your “ImportDirectory” with .R4 appended to it – so for example ./vonk-import.R4.

#### ImportedDirectory

This directory will contain the read history in the .vonk-import-history.json file. Please note, that this information is stored directly in the administration database when running on SQLite.

Note that in json you either use forward slashes (/) or double backward slashes (\\) as path separators.

### 5.4.6 Load Conformance Resources from simplifier.net

You are encouraged to manage and publish your profiles and related Conformance Resources on [simplifier.net](https://simplifier.net). If you do that, you can have Firely Server read those. You configure this in the *Firely Server settings*:

```
"AdministrationImportOptions": {
  "SimplifierProjects": [
    {
      "Uri": "FHIR endpoint for retrieving StructureDefinitions",
      "UserName": "UserName for retrieving the StructureDefinitions",
      "Password": "Password for the above user name",
      "BatchSize": "<number of resources imported at once, optional - default is 20>"
    }
  ],
}
```

#### Uri

must point to a Simplifier project endpoint, see below on how to get this

#### UserName

your username, if you access a private Simplifier project

#### Password

password with the username

#### BatchSize

you normally don't need to change this parameter

You can load from multiple Simplifier projects by adding them to the list. The environment variable version of this is:

```
VONK_Administration:SimplifierProjects:0:Uri=<FHIR endpoint for retrieving
↳StructureDefinitions>
```

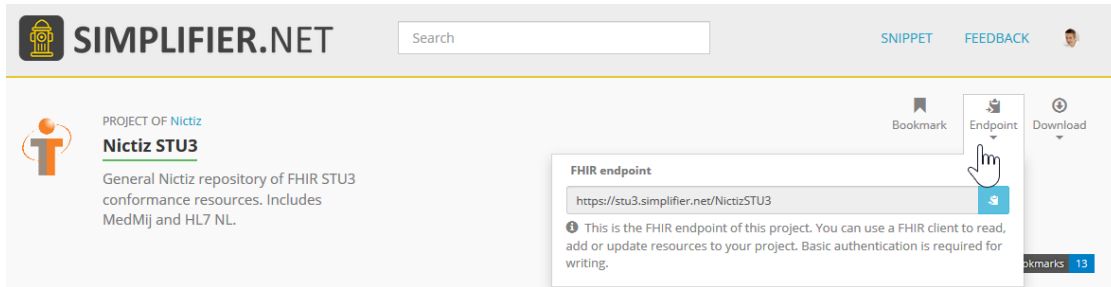
Vonk automatically finds the FHIR version for each project and imports it only for the matching FHIR version.

### Get a FHIR endpoint for a Simplifier project

Open the project of your choice on <https://simplifier.net>. There are two limitations:

1. You must have access to the project (so either public or private but accessible to you)
2. The project must be STU3

Then on the overview page of the project click 'Endpoint' and copy the value you see there:



By default the endpoint is `https://stu3.simplifier.net/<projectname>`

## 5.4.7 Load Conformance Resources on demand

It can be useful to reload the profiles, e.g. after you have finalized changes in your project. Therefore you can instruct Firely Server to actually load the profiles from the source(s) with a separate command:

```
POST http(s)://<firely-server-endpoint>/administration/importResources
```

The operation will return an `OperationOutcome` resource, containing details about the number of resources created and updated, as well as any errors that occurred. Please note that this will also respect the history of already read files, and not read them again.

## 5.4.8 Manage Conformance Resources with the Administration API

The *Firely Server Administration API* has a FHIR interface included, on the `https://<firely-server-endpoint>/administration` endpoint. On this endpoint you can do most of the FHIR interactions (create, read, update, delete, search) on these resourcetypes:

- SearchParameter
- StructureDefinition
- ValueSet
- CodeSystem
- CompartmentDefinition

If you are *not permitted* to access the endpoint for the resource you want to manage (e.g. `<firely-server-endpoint>/administration/StructureDefinition`), Firely Server will return status-code 403.

**Note:** You can also do the same interactions on the same resourcetypes on the normal (or 'data') FHIR endpoint `https://<firely-server-endpoint>`. This will only result in storing, updating or deleting the resource. But it will not have any effect on the way Firely Server operates.

---

### Example

To add a StructureDefinition to Firely Server

```
POST <firely-server-endpoint>/administration/StructureDefinition
```

- In the body provide the StructureDefinition that you want to add.
- The Content-Type header must match the format of the body (application/fhir+json or application/fhir+xml)

If you prefer to assign your own logical id to e.g. StructureDefinition 'MyPatient', you can use an update:

```
PUT <firely-server-endpoint>/administration/StructureDefinition/MyPatient
```

## 5.5 Validating incoming resources

You can have Firely Server validate all resources that are sent in for create or update. The setting to do that is like this:

```
"Validation": {
  "Parsing": "Permissive", // Permissive / Strict
  "Level": "Off", // Off / Core / Full
  "AllowedProfiles":
  [
    "http://hl7.org/fhir/StructureDefinition/daf-patient",
    "http://hl7.org/fhir/StructureDefinition/daf-allergyintolerance"
  ]
},
```

### 5.5.1 Parsing

Every incoming resource - xml or json - has to be syntactically correct. That is not configurable.

Beyond that, you can choose between Permissive or Strict parsing. Permissive allows for:

- empty elements (not having any value, child elements or extensions)
- the `fhir_comments` element
- errors in the xhtml of the Narrative
- **json specific:**
  - array with a single value instead of just a value, or vice versa (json specific)
- **xml specific:**
  - repeating elements interleaved with other elements
  - elements out of order

- mis-representation (element instead of attribute etc.)

### 5.5.2 Validation

You can choose the level of validation:

- Off: no validation is performed.
- Core: the resource is validated against the core definition of the resourcetype.
- Full: the resource is validated against the core definition of the resourcetype and against any profiles it its `meta.profile` element.

### 5.5.3 Allow for specific profiles

To enable this feature, set `Level` to `Full`.

If you leave the list of `AllowedProfiles` empty, any resource will be allowed (provided it passes the validations set in `Parsing` and `Level`).

When you add canonical urls of `StructureDefinitions` to this list, Firely Server will:

- check whether the incoming resource has any of these profiles listed in its `meta.profile` element
- validate the resource against the profiles listed in its `meta.profile` element.

So in the example above, Firely Server will only allow resources that conform to either the DAF Patient profile or the DAF AllergyIntolerance profile.

Note that the resource has to declare conformance to the profile in its `meta.profile` element. Firely Server will *not* try to validate a resource against all the `Validation.AllowedProfiles` to see whether the resource conforms to any of them, only those that the resource claims conformance to.

### 5.5.4 Upgrading from < 2.0

Previous to version 2.0.0, the setting was `"ValidateIncomingResources": "true" // or false`. The corresponding settings since version 2.0.0 are:

- true / no `AllowedProfiles` => `"Level": "Core"`
- true / with `AllowedProfiles` => `"Level": "Full"`, keep the `AllowedProfiles` as it is.
- false => `"Level": "Off"`

## 5.6 Using the In-Memory storage

- Navigate to your Firely Server working directory
- Changing a setting means overriding it as described in *Changing the settings*.
- Find the `Repository` setting:

```
"Repository": "Memory",
```

- If it is not already set to `Memory`, do so now.

- You can set `SimulateTransactions` to “true” if you want to experiment with [FHIR transactions](#). The In-Memory implementation does not support real transactions, so in case of an error already processed entries will NOT be rolled back:

```
"MemoryOptions": {  
  "SimulateTransactions": "true"  
},
```

### 5.6.1 Using the In-Memory storage for the Administration API database

This works the same as with the normal Firely Server database, except that you put the settings within the Administration section

E.g.:

```
"Administration": {  
  "Repository": "Memory",  
  "MemoryOptions": {  
    "SimulateTransactions": "false"  
  }  
}
```

**Warning:** Using the In-Memory storage for the administration layer will cause Firely Server to load the specification files on each startup. This takes several minutes, and Firely Server will respond with a ‘423 - Locked’ error to all requests during that time. As of version 0.7.1 we have implemented support for SQLite, which we recommend to use instead of the In-Memory storage. See [Prefer SQLite for Firely Server Administration](#) for more information.

## 5.7 Using MongoDB

We assume you already have MongoDB installed. If not, please refer to the [MongoDB download](#) pages.

Firely Server can work with MongoDB 4.0 and higher. Since Firely Server (Vonk) version 3.7.0 Firely Server uses the MongoDB Aggregation Framework heavily, and you are advised to upgrade to MongoDB 4.4 (or newer). In this version issue [SERVER-7568](https://jira.mongodb.org/browse/SERVER-7568) <<https://jira.mongodb.org/browse/SERVER-7568>> is solved, so the most selective index is used more often.

---

**Note:** Using the correct Read and Write Concern for your MongoDB replica set is very important:

- Write Concern:
    1. For more information, check the MongoDB manual about [Write Concern](#)
    2. For PSA (Primary-Secondary-Arbiter) replica sets you would want the Write Concern to be “w=1”
    3. For PSS (Primary-Secondary-Secondary) replica sets you would want the Write Concern to be “w=majority”
  - Read Concern:
    1. For more information, check the MongoDB manual about [Read Concern](#)
    2. Default Read Concern is *local*.
-



- Navigate to your Firely Server working directory
- Changing a setting means overriding it as described in *Changing the settings*.
- Find the Repository setting:

```
"Repository": "SQLite",
```

- Change the setting to MongoDB
- If you have your own database in MongoDB already, change the `MongoDbOptions` to reflect your settings:

```
"MongoDbOptions": {
  "ConnectionString": "mongodb://localhost/vonkdata",
  "EntryCollection": "vonkentries"
},
```

- If MongoDB does not have a database and/or collection by this name, Firely Server will create it for you.
- Find the section called `PipelineOptions`. Make sure it contains the MongoDB repository in the root path for Firely Server Data:

```
"PipelineOptions" :
{
  "Branches" : [
    "/" : {
      "Include" : [
        "Vonk.Repository.MongoDb.MongoDbVonkConfiguration"
        //...
      ]
    }
  ]
}
```

### 5.7.1 Using MongoDB for the Administration API database

Although we encourage you to use *SQLite for Firely Server Administration*, you can still use MongoDB for Firely Server Administration as well.

This works the same as with the normal Firely Server database, except that you:

- put the settings within the `Administration` section
- provide a different `ConnectionString` and/or `EntryCollection`, e.g.:

```
"Administration": {
  "Repository": "MongoDB",
  "MongoDbOptions": {
    "ConnectionString": "mongodb://localhost/vonkadmin",
    "EntryCollection": "vonkadmin"
  }
}
```

- Find the section called `PipelineOptions`. Make sure it contains the MongoDB repository in the administration path for Firely Server Administration:

```
"PipelineOptions" :
{
  "Branches" : [
    "/administration" : {
      "Include" : [
        "Vonk.Repository.MongoDb.MongoDbAdministrationConfiguration"
        //...
      ]
    }
  ]
}
```

**Attention:** For MongoDB it is essential to retain the `.vonk-import-history.json` file. Please read [Retain the import history](#) for details.

## 5.7.2 MongoDB Transactions

---

**Note:** When utilizing MongoDB transactions we strongly advise to use MongoDB v4.2 or higher.

---

In Firely Server versions prior to v4.9.0 transactions were simulated for development and test purposes. From Firely Server v4.9.0 and onwards transactions using MongoDB are now fully supported.

With MongoDB transactions, there are a few things to consider:

1. MongoDB supports transactions only for *Replica Sets* and *Sharded Clusters*. If you are running Firely Server on a MongoDB standalone instance you still will be able to upload a transaction bundle, but it will not be processed within a transaction. I.e.: if an exception occurs with a resource during processing the bundle, any previous resources will have been persisted to the database and not rolled back.
2. Firely Server currently uses transactions in the following cases:
  1. When uploading a transaction bundle.
  2. When performing a conditional delete that targets more than one resource.
3. MongoDB transactions in Firely Server always use Read Concern “*snapshot*” and Write Concern “*majority*”.
4. MongoDB imposes a transaction runtime limit of *60s*. For self-hosted MongoDB instances you can modify this limit using “*transactionLifetimeLimitSeconds*”. However, for MongoDB Atlas deployments this limit cannot be changed.
5. Although MongoDB transactions are supported as early as v4.0, please be aware of the following issue. In MongoDB v4.0 all write operations are contained in a single oplog entry. The oplog entry for the transaction must be within the BSON document size limit of 16MB. For v4.2+ every write operation gets its own oplog entry. This removes the 16MB total size limit for a transaction imposed by the single oplog entry for all its write operations. Note that each single oplog entry still has a limit of 16 MB. We highly recommend in using MongoDB v4.2 or higher when using transactions.
6. Please read the official MongoDB documentation for production considerations when using transactions: [MongoDb manual](#)

### 5.7.3 Tips and hints for using MongoDB for Firely Server

1. If searches and/or creates and updates are excessively slow, you may be limited by the IOPS on your MongoDB deployment (e.g. MongoDB Atlas). Try upgrading it and check the timings again.
2. If for any reason you would like to see how Firely Server is interacting with MongoDB, make the following adjustments to the [Log settings](#):
  1. In the section `Serilog.MinimumLevel.Override` add `"Vonk.Repository.DocumentDb": "Verbose"`. Add it before any broader namespaces like `Vonk`.
  2. In the section on the File sink, change the `restrictedToMinimumLevel` to `Verbose`.
3. **With regards to Firely Server version and MongoDB version:**
  1. If you are on a Firely Server (Vonk) version < v3.6, you can keep using MongoDB v4.0 or higher.
  2. If you are on Firely Server (Vonk) v3.6 or higher and are unable to migrate to MongoDB 4.4 (relatively soon), please contact us if you need assistance.

### 5.7.4 MongoDB disk space requirements

A MongoDB database has no single size, but several characteristics that tell us something about the size of the data:

#### **nr of documents**

This is not indicative of any storage size, because documents can be sized very differently.

#### **document size**

The size in bytes of all the documents in a database, uncompressed.

#### **index size**

The size in bytes of all the indexes in a database, uncompressed.

#### **storage size**

The actual space used for storage of the documents and indexes. Since the storage engine of MongoDB (WiredTiger) compresses data by default, this is often substantially less than the document and index sizes combined.

#### **disk space used**

The size of the database file on disk. Because data can become fragmented, this is larger than the storage size.

The sizes that are relevant for estimating the disk space requirements are the storage size and the disk space used. At Firely we host a test instance on a MongoDB Atlas M40 cluster, with 1 primary and 2 replicas. Based on that we can calculate these sizes:

1. storage size: ~ 1 GB per 1.000.000 (1 mln) resources
2. disk size: ~ 2 times the storage size, so 2 GB per 1 mln resources, *per replica*, allowing for a fragmentation ratio up to 50%
3. buffer: 20% of disk size (see below)

Storage size may differ based on the size of the resources you host. These estimations are based on Synthea patient records. But e.g. `ExplanationOfBenefit` resources are typically much larger than the average `Observation` resource.

Fragmentation in MongoDB occurs when data is deleted. New data is appended at the end of the data file. Firely Server will delete data upon update, delete and \$erase operations. This means that you should account for more or less fragmentation based on how many of these operations you expect will be performed on your Firely Server instance. If fragmentation gets too high, you can [compact](#) the collection with resources (by default named `vonkentries`).

On top of the requirements for storing the resources and indexes, we allow MongoDB to use the disk as an overflow buffer for larger-than-memory operations (like sorting a very large resultset). Please reserve about 20% extra disk space for that.

We recommend to monitor the health of your MongoDB cluster actively to avoid disk space issues.

## 5.8 Using SQL server

There are two ways to create the Firely Server database on a SQL Server instance: Have Firely Server create it for you entirely or create an empty database and users yourself and have Firely Server create the schema (tables etc.).

In both cases:

- Prepare an instance of SQL Server 2012 or newer. Any edition - including SQL Server Express - will do. The instance will have a servername and possibly an instancename: `server/instance`.
- Changing a setting means overriding it as described in *Changing the settings*.
- Find the `Repository` setting:

```
"Repository": "SQLite",
```

- Change the setting to SQL
- Find the section called `SqlDbOptions`. It has these values by default:

```
"SqlDbOptions": {
  "ConnectionString": "connectionstring to your Firely Server SQL Server database,
↳ (SQL2012 or newer); Set MultipleActiveResultSets=True",
  "SchemaName": "vonk",
  "AutoUpdateDatabase": true,
  "MigrationTimeout": 1800 // in seconds
  //"AutoUpdateConnectionString" : "set this to the same database as
↳ 'ConnectionString' but with credentials that can alter the database. If not set,
↳ defaults to the value of 'ConnectionString'"
},
```

- Find the section called `PipelineOptions`. Make sure it contains the SQL repository in the root path for Firely Server Data. For Firely Server versions older than 4.3.0 use `Vonk.Repository.Sql.SqlVonkConfiguration`. For Firely Server v4.3.0 and above use `Vonk.Repository.Sql.Raw.KSearchConfiguration`:

```
"PipelineOptions" :
{
  "Branches" : [
    "/" : {
      "Include" : [
        //"Vonk.Repository.Sql.SqlVonkConfiguration", // use this for FS,
↳ versions < v4.3.0
        "Vonk.Repository.Sql.Raw.KSearchConfiguration", // use this for FS,
↳ versions >= v4.3.0
        //...
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```
    ]
  }
```

- The site [connectionstrings.com](https://connectionstrings.com) is useful for determining the correct connectionstring for your environment.
- If you will only use Windows Accounts, you can use the (default) Authentication Mode, which is Windows Authentication Mode. But if you also want to use SQL Server accounts, you have to run it in Mixed Mode. Refer to [Authentication in SQL Server](#) for more information.
- Although we encourage you to use *SQLite for Firely Server Administration*, you can still use SQL Server for Firely Server Administration as well:

```
"Administration": {
  "Repository": "SQL",
  "SqlDbOptions": {
    "ConnectionString": "Integrated Security=SSPI;Persist Security Info=False;
↪Initial Catalog=VonkAdmin;Data Source=Server\\Instance;
↪MultipleActiveResultSets=true",
    "SchemaName": "vonk",
    "AutoUpdateDatabase": true,
    "MigrationTimeout": 1800 // in seconds
  }
},
//...
"PipelineOptions" :
{
  "Branches" : [
    "/administration" : {
      "Include" : [
        //"Vonk.Repository.Sqlite.SqliteTaskConfiguration", // use this for ↪
↪FS versions < v4.3.0
        "Vonk.Repository.Sql.Raw.KAdminSearchConfiguration", // use this ↪
↪for FS versions >= v4.3.0
        //...
      ]
    }
  ]
}
```

### 5.8.1 Have Firely Server create your database

This option is mainly for experimentation as it effectively requires sysadmin privileges for the connecting user.

- Prepare a login on SQL Server with the following role:
  - sysadmin
- Set the SqlDbOptions for the Firely Server database as follows (the values are example values for connecting with your own Windows login):

```
"SqlDbOptions": {
  "ConnectionString": "Integrated Security=SSPI;Persist Security Info=False;
↪Initial Catalog=VonkData;Data Source=Server\\Instance;MultipleActiveResultSets=true
```

(continues on next page)

(continued from previous page)

```

    ↪",
    "SchemaName": "vonk",
    "AutoUpdateDatabase": true,
    "MigrationTimeout": 1800 // in seconds
  },

```

- Set the `SqlDbOptions` under `Administration` for the `Administration` database likewise:

```

"Administration": {
  "Repository": "SQL",
  "SqlDbOptions": {
    "ConnectionString": "Integrated Security=SSPI;Persist Security Info=False;
    ↪Initial Catalog=VonkAdmin;Data Source=Server\Instance;
    ↪MultipleActiveResultSets=true",
    "SchemaName": "vonk",
    "AutoUpdateDatabase": true,
    "MigrationTimeout": 1800 // in seconds
  }
}

```

- You don't need to set `AutoUpdateConnectionString` since the `ConnectionString` will already have enough permissions.
- Start Firely Server. It will display in its log that it applied pending migrations. After that the database is created and set up with the correct schema.
- If an upgrade to a new version of Firely requires a migration then a SQL time out might occur, halting the upgrade and resulting in a rollback of the migration. The duration of the SQL time out for migrations can be controlled with `MigrationTimeout`. The default value is 1800 seconds (30 min).

**Attention:** For SQL Server it is essential to retain the `.vonk-import-history.json` file. Please read [Retain the import history](#) for details.

### 5.8.2 Create a database and users by script, and have Firely Server create the schema

- Log into SQL Server as the Administrator user.
- From the working directory open `data01-CreateDatabases.sql`
- In SQL Server Management Studio, in the menu select Query|SQLCMD Mode.
- In the script uncomment and adjust the variable names `dbName` and `AdminDbName` as well as any other variables to your own liking.
- Run the script to create both the Firely Server database and the Administration API database.
- From the working directory open `data02-CreateDBUser.sql`
- In SQL Server Management Studio, in the menu select Query|SQLCMD Mode.
- In the script uncomment and adjust the variables at the top names to your own liking.
- Run the script to create two users, one with access to the Firely Server database, the other with access to the Administration database. This script grants the database role `db_ddladmin` to both users, to enable the `AutoUp-`

dateDatabase feature. Refer to *Overview of permissions* for an overview of necessary authorization for different features.

- Set the SqlDbOptions for the Firely Server database as follows:

```
"SqlDbOptions": {
  "ConnectionString": "User Id=<dbUserName>;Password=<dbPassword>;Initial Catalog=
↪<DataDbName>;Data Source=server\\instance;MultipleActiveResultSets=True",
  "SchemaName": "vonk",
  "AutoUpdateDatabase": "true"
}
```

- If you have set up a different user for running the AutoUpdateDatabase feature, you can provide that:

```
"SqlDbOptions": {
  "ConnectionString": "User Id=<dbUserName>;Password=<dbPassword>;Initial Catalog=
↪<DataDbName>;Data Source=server\\instance;MultipleActiveResultSets=True",
  "SchemaName": "vonk",
  "AutoUpdateDatabase": "true"
  "AutoUpdateConnectionString": "User Id=<updateUserName>;Password=
↪<updatePassword>;Initial Catalog=<DataDbName>;Data Source=server\\instance;
↪MultipleActiveResultSets=True",
}
```

- Set the SqlDbOptions under Administration for the Administration database likewise:

```
"Administration" : {
  "Repository": "SQL",
  "SqlDbOptions": {
    "ConnectionString": "User Id=<AdminDbUserName>;Password=<AdminDbPassword>;
↪Initial Catalog=<AdminDbName>;Data Source=server\\instance;
↪MultipleActiveResultSets=True",
    "SchemaName": "vonk",
    "AutoUpdateDatabase": "true"
  }
}
```

- For the administration you can also provide different credentials for performing the auto update:

```
"Administration" : {
  "Repository": "SQL",
  "SqlDbOptions": {
    "ConnectionString": "User Id=<AdminDUserName>;Password=<AdminDbPassword>;
↪Initial Catalog=<AdminDbName>;Data Source=server\\instance;
↪MultipleActiveResultSets=True",
    "SchemaName": "vonk",
    "AutoUpdateDatabase": "true"
    "AutoUpdateConnectionString": "User Id=<updateAdminUserName>;Password=
↪<updateAdminPassword>;Initial Catalog=<AdminDbName>;Data Source=server\\instance;
↪MultipleActiveResultSets=True",
  }
}
```

### 5.8.3 Run schema migrations manually

It is possible to disable automatic migrations for the database, by setting `AutoUpdateDatabase` to false. Regardless, with a new version, the schema must be upgraded for the application to run. With auto-updating disabled, an error will halt the startup. The error contains the version of the current schema and the required schema version of the server. You must use the SQL script(s) in the `/sqlserver` folder to update your database to the required version.

### 5.8.4 Overview of permissions

This table lists the permissions needed to perform specific actions on the SQL database. Recommended roles are listed in the third column. Note that you can create your own database roles with the required permissions in order to execute finely granulated control over permissions.

Table 1: Permissions and Roles

Action	Required SQL Permission	Recommended SQL Role	Notes
Create Database	'Create any database'	sysadmin	
AutoUpdateDatabase feature including application of pending migrations	'Alter Trace' + 'Alter any database'	db_ddladmin	Applies to the normal Firely Server database and its administration database
AutoUpdateDatabase feature excluding application of pending migrations	'Alter Trace'	db_ddladmin	
Read resources		db_datareader	
Write resources		db_datawriter	
Execute ResetDb feature		db_ddladmin	Only applies to the normal Firely Server database

If the `AutoUpdate` feature is enabled, we recommend creating two users in order to achieve maximum security while keeping the configuration simple:

- User 1 with the 'db\_datareader' and 'db\_datawriter' and 'db\_ddladmin' roles (This user should be used in the `AutoUpdateConnectionString`)
- User 2 with the 'db\_datareader' and 'db\_datawriter' roles as well as 'Alter Trace' permissions for day to day usage (This user should be used in the `ConnectionString`)

## 5.9 Using SQLite

SQLite is a file based database engine. The engine itself does not run separately on the server, but in-process in the application, Firely Server in this case.

For more background on SQLite please refer to the [SQLite documentation](#).

SQLite is the default configuration of Firely Server. For the Administration database there is little reason to change this. For the actual runtime data, (the 'Firely Server database') itself, you may run into limitations of SQLite if you put it through its paces. You may find one of the other repositories a better fit then. You can safely use different storage engines for Firely Server Data and Firely Server Administration.



### 5.9.1 Prefer SQLite for Firely Server Administration

Until Firely Server (Vonk) version 0.7.0 you could use any of the storage engines for both Firely Server Data and Firely Server Administration. Starting with Firely Server (Vonk) 0.7.1 you are encouraged to use SQLite for Firely Server Administration. Over time we will deprecate support for running Firely Server Administration on the SQL Server, MongoDB and Memory storage engines. For Firely Server Data you can of course still use the storage engine of your preference.

Firely Server Administration poses very limited stress on its storage engine, therefore SQLite is adequate. And it provides several advantages:

- Runs out of the box: SQLite requires no installation of a database engine, but still provides durable storage (unlike the Memory storage). Thus, you don't need to setup anything to run Firely Server Administration. And you can download the Firely Server binaries and run them without any further configuration.
- Flexible on updates: Many of the features that we will add to Firely Server require changes to the schema of the Administration database. By only supporting SQLite for this, we can provide these features to you more quickly.
- Readymade database: In the other storage engines, the conformance resources from the specification had to be *imported* before Firely Server could start. This would take a couple of minutes. Because SQLite is file based, we can run the import process for you and provide you with a readymade Administration database.
- Runs with Facades: perhaps the most important feature. If you build a Firely Server Facade, the facade will not provide support for hosting conformance resources. With Firely Server Administration on SQLite the facade has its own storage and you can use Firely Server Administration out of the box. This enables e.g. validation against your custom resources (that can be imported from your Simplifier project), subscriptions, and other use cases.

### 5.9.2 Settings for using SQLite for Firely Server Data

- Changing a setting means overriding it as described in *Changing the settings*.
- Find the `Repository` setting and set it to SQLite if it not already set to that:

```
"Repository": "SQLite",
```

- Find the section called `SQLiteDbOptions`. It has these values by default:

```
"SQLiteDbOptions": {
  "ConnectionString": "Data Source=./data/vonkdata.db",
  "AutoUpdateDatabase": true
},
```

Firely Server will create the database *file*, but please make sure the *directory* already exists.

- Find the section called `PipelineOptions`. Make sure it contains the SQLite repository in the root path:

```
"PipelineOptions" :
{
  "Branches" : [
    "/" : {
      "Include" : [
        "Vonk.Repository.SQLite.SQLiteVonkConfiguration"
        //...
      ]
    },
    //...
```

(continues on next page)

(continued from previous page)

```
]
}
```

### 5.9.3 Settings for using SQLite for Firely Server Administration

- Set the `SqlDbOptions` under `Administration` for the `Administration` database similar to those above:

```
"Administration" : {
  "Repository": "SQLite",
  "SqliteDbOptions": {
    "ConnectionString": "Data Source=./data/vonkadmin.db",
    "AutoUpdateDatabase": "true"
  }
}
```

Firely Server will create the database *file*, but please make sure the *directory* already exists.

- Find the section called `PipelineOptions`. Make sure it contains the `SQLite` repository in the administration path:

```
"PipelineOptions" :
{
  "Branches" : [
    "/" : {
      //...
    },
    "/administration" : {
      "Include" : [
        "Vonk.Repository.SQLite.SqliteAdministrationConfiguration"
        //...
      ]
    }
  ]
}
```

### 5.9.4 Administration import history in SQLite

When Firely Server *imports Conformance resources*, it keeps record of what is has imported. Unlike the SQL Server and MongoDB engines, the SQLite storage engine does *not* use the `.vonk-import-history.json` file for that. Instead, in SQLite the import history is stored within the Administration database itself.

## 5.10 Configure http and https

You can enable http and/or https and adjust the port numbers for them in *Firely Server settings*.

### 5.10.1 Changing the port number

By default Firely Server will run on port 4080 of your system. You can change the port setting by overriding it as described in *Changing the settings*:

- Navigate to your Firely Server working directory
- Find this setting:

```
"Hosting": {
  "HttpPort": 4080
}
```

- Change the number to the port number you want

### 5.10.2 Changing from http to https

If you need your server to run on https instead of http, follow these steps:

- Navigate to the location where you extracted the Firely Server files.
- Find these settings:

```
"Hosting": {
  "HttpPort": 4080,
  "HttpsPort": 4081, // Enable this to use https
  "CertificateFile": "<your-certificate-file>.pfx", //Relevant when
↪HttpsPort is present
  "CertificatePassword" : "<cert-pass>" // Relevant when HttpsPort is
↪present
},
```

- Uncomment the lines for HttpsPort, CertificateFile and CertificatePassword.
- Set the HttpsPort to the port of your liking (standard https port is 443)
- Set CertificateFile to the location of the `.pfx` file that contains the certificate for your site
- Set CertificatePassword to the password for the certificate file.

**Note:** We recommend setting this value as an environment variable for security reasons:

```
VONK_Hosting:CertificatePassword=<password>
```

To set this:

- In Powershell run: `> $env:VONK_Hosting:CertificatePassword="my_password"` where *my\_password* is the password for the *.pfx* file
- or go to your *System*, open the *Advanced system settings* → *Environment variables* and create a new variable with the name `VONK_Hosting:CertificatePassword` and the value set to your password
- You can choose to comment-out the `HttpPort` setting, so Firely Server will no longer be available through unsecured http.

## 5.11 Cross Origin Resource Sharing (CORS)

CORS is enabled in Firely Server. Since Firely Server provides an API that is expected to be consumed by applications from different domains, CORS is enabled for any origin. Currently there is no setting to control this behaviour.

## 5.12 Log settings

Firely Server uses [Serilog](#) for logging. The logging settings are controlled in json configuration files called `logsettings(.*) .json`. The files are read in a hierarchy, exactly like the [appsettings files](#) are. Firely Server comes with default settings in `logsettings.default.json`. You can adjust the way Firely Server logs its information by overriding these default settings in `logsettings.json` or `logsettings.instance.json`. You need to create this file yourself.

Alternatively you can control *Log settings with Environment Variables*.

Firely Server by default does not log any Patient Health Information data, regardless of the level of the log. The only PHI part that can be included in the log is the User Name, when running with Smart authorization (so the user is identified) and including this part in the outputTemplate (see below).

### 5.12.1 Changing the log event level

Serilog defines several levels of log events. From low to high, these are `Verbose`, `Debug`, `Information`, `Warning`, `Error` and `Fatal`. You can set the minimum level you want to log, meaning that events for that level or higher will be logged. By default, Firely Server uses `Error` as the minimum level of recording information.

To change the level of logging, follow these steps:

- Find this setting:

```
"MinimumLevel": {  
  "Default": "Error",  
},
```

- Change the setting for `Default` from `Error` to the level you need, from the choice of `Verbose`, `Debug`, `Information`, `Warning`, `Error` and `Fatal`.

You can deviate from the default minimum level for specific namespaces. You do this by specifying the namespace and the log event level you would like for this namespace, for example:

```
"MinimumLevel": {  
  "Default": "Error",  
  "Override": {
```

(continues on next page)

(continued from previous page)

```

        "Vonk": "Warning"
    },
},

```

Some additional namespaces you might want to log are:

- `Vonk.Configuration` to log configuration information on startup
- `Vonk.Core.Licensing` to show license information in your logs
- `Vonk.Repository.Sql.Raw` to log SQL repository events for Firely Server v4.3.0 and above
- `Vonk.Repository.Sql` to log SQL repository events for Firely Server v4.2.0 and below
- `Vonk.Repository.Document.Db` to log MongoDB repository events
- `Vonk.Repository.Memory` to log memory database repository events
- `Vonk.Core.Repository.EntryIndexerContext`, set it to "Error" if you have excessive warnings about indexing (mostly when importing *Synthea* <<https://synthea.mitre.org/downloads>> data)
- `Microsoft` to log events from the Microsoft libraries
- `Microsoft.AspNetCore.Diagnostics` to report request handling times
- `System` to log events from the System libraries

Please note that the namespaces are evaluated in order from top to bottom, so more generic 'catch all' namespaces should be at the bottom of the list. So this will log events on `Vonk.Repository.Sql.Raw` on `Information` level:

```

"MinimumLevel": {
    "Default": "Error",
    "Override": {
        "Vonk.Repository.Sql.Raw": "Information",
        "Vonk": "Warning"
    }
},

```

But in this (purposefully incorrect) example the `Warning` level on the `Vonk` namespace will override the `Information` level on the `Vonk.Repository.Sql.Raw` namespace:

```

"MinimumLevel": {
    "Default": "Error",
    "Override": {
        "Vonk": "Warning",
        "Vonk.Repository.Sql.Raw": "Information"
    }
},

```

### 5.12.2 Changing the sink

Another setting you can adjust is `WriteTo`. This tells Serilog which sink(s) to log to. Serilog provides several sinks, and for Firely Server you can use `Console`, `File`, `ApplicationInsights` and `Seq`. All of which can be wrapped in an `Async` sink to avoid blocking Firely Server when waiting for the sink to process the log statements.

#### Console

The Console sink will write to your shell.

- Find the `WriteTo` setting:

```
"WriteTo": [
  {
    "Name": "Async",
    "Args": {
      "configure": [
        {
          "Name": "Console",
          "Args": {
            "restrictedToMinimumLevel": "Information",
            "outputTemplate": "{Timestamp:yyyy-MM-dd,
↪HH:mm:ss.fff zzz} {UserId} {Username} [{Level}] [ReqId: {RequestId}] {Message}
↪{NewLine}{Exception}"
          }
        }
      ]
    }
  },
  {
    //Settings for other sinks
  }
]
```

The Console is notoriously slow at processing log statements, so it is recommended to limit the number of statements for this sink. Use the `restrictedToMinimumLevel` to do so. Also, if you are on Windows, the Powershell command window appears to be faster than the traditional Command Line window.

Settings for the Console sink:

- `outputTemplate`: What information will be in each log line. Besides regular text you can use placeholders for information from the log statement:
  - `{Timestamp:yyyy-MM-dd HH:mm:ss.fff zzz}`: When this was logged, with formatting
  - `{UserId}`: Technical id of the logged in user - if applicable
  - `{Username}`: Name of the logged in user - if applicable
  - `{Application}`: Name of the application (in case other applications are logging to the same sink). Is set to Vonk at the bottom of the logsettings file
  - `{Level}`: Level of the log, see the values in [Changing the log event level](#)
  - `{MachineName}`: Name of the machine hosting the Firely Server instance. Especially useful when running multiple instances all logging to the same file.
  - `{RequestId}`: Unique id of the web request, useful to correlate log statements
  - `{Message}`: Actual message being logged

- {Exception}: If an error is logged, Firely Server may include the original exception. That is then formatted here.
- {SourceContext}: The class from which the log statement originated (this is usually not needed by end users).
- {NewLine}: Well, eh, continue on the next line
- restrictedToMinimumLevel: Only log messages from this level up are sent to this sink.

## File

The File sink will write to a file, possibly rolling it by interval or size.

- Find the WriteTo setting:

```
"WriteTo": [
  {
    {
      //Settings for Console
    }
  },
  {
    "Name": "Async",
    "Args": {
      "configure": [
        {
          "Name": "File",
          "Args": {
            "path": "%temp%/vonk.log",
            "rollingInterval": "Day",
            "fileSizeLimitBytes": "",
            "retainedFileCountLimit": "7",
            "outputTemplate": "{Timestamp:yyyy-MM-dd HH:mm:ss.
↪fff zzz} {UserId} {Username} [{Application}] [{Level}] [Machine: {MachineName}] ↪
↪[ReqId: {RequestId}] {Message}{NewLine}{Exception}",
            "restrictedToMinimumLevel": "Verbose"
          }
        }
      ]
    }
  },
  {
    //Settings for Azure ApplicationInsights
  }
]
```

- Under File, change the location of the logfiles by editing the value for path. For example:

```
{
  "Name": "RollingFile",
  "Args": {
    "path": "c:/logfiles/vonk.log"
  }
},
```

Other values that you can set for the File log are:

- `rollingInterval`: When this interval expires, the log system will start a new file. The start datetime of each interval is added to the filename. Valid values are `Infinite`, `Year`, `Month`, `Day`, `Hour`, `Minute`.
- `fileSizeLimitBytes`: Limit the size of the log file, which is 1GB by default. When it is full, the log system will start a new file.
- `retainedFileCountLimit`: If more than this number of log files is written, the oldest will be deleted. Default value is 31. Explicitly setting it to an empty value means files are never deleted.
- `outputTemplate`: as described for *Console*.
- `restrictedToMinimumLevel`: as described for *Console*.

### Application Insights

Firely Server can also log to Azure Application Insights (“Application Insights Telemetry”). What you need to do:

1. Create an Application Insights instance on Azure.
2. Get the `InstrumentationKey` from the Properties blade of this instance.
3. Add the correct sink to the `logsettings.json`:

```
"WriteTo": [  
  {  
    "Name": "ApplicationInsightsTraces",  
    "Args": {  
      "instrumentationKey": "<the key you copied in step 2>",  
      "restrictedToMinimumLevel": "Verbose" //Or a higher level  
    }  
  },  
],
```

4. This also enables Dependency Tracking for access to your database. This works for both SQL Server and MongoDB. And for the log sent to *Seq* if you enabled that.
5. If you set the level for Application Insights to `Verbose`, and combine that with *Database details*, you get all the database commands right into Application Insights.

### Seq

*Seq* is a web interface to easily inspect structured logs.

For the *Seq* sink, you can also specify arguments. One of them is the server URL for your *Seq* server:

```
"WriteTo": [  
  {  
    "Name": "Seq",  
    "Args": { "serverUrl": "http://localhost:5341" }  
  }  
]
```

- Change `serverUrl` to the URL of your *Seq* server
- `restrictedToMinimumLevel`: as described for *Console*.



### 5.12.3 Database details

Whether you use MongoDB or SQL Server, you can have Firely Server log in detail what happens towards your database. Just set the appropriate loglevel to 'Verbose':

```
"MinimumLevel": {
  "Default": "Error",
  "Override": {
    "Vonk.Repository": "Verbose"
  }
},
```

You can also control the logsettings for the different repositories more finely granulated:

```
"MinimumLevel": {
  "Default": "Error",
  "Override": {
    // (for versions before FS 4.6.0)
    "Vonk.Repository.Sql": "Verbose",
    // OR (for FS 4.6.0 or later AND if Sql.Raw is enabled)
    "Vonk.Repository.Sql.Raw": "Verbose",
    // OR (for MongoDB)
    "Vonk.Repository.Document.Db": "Verbose",
    // OR (for SQLite)
    "Vonk.Repository": "Verbose",
    "Microsoft.EntityFrameworkCore": "Verbose"
  }
},
```

Remember to adjust your sink settings so that "restrictedToMinimumLevel": "Verbose" is set. If you do so you probably don't want all this detail in your console sink, so you can limit the level for that, see [Console](#) above.

### SQL query parameter logging

It might be useful to log SQL queries that Firely Server executes against your database. You can get even more insights into what is happening when SQL query parameter values also get logged. However, this cannot be enabled by default due to data privacy concerns.

You can enable SQL query parameter values logging by setting the LogSqlQueryParameterValue to true for the corresponding database in your appsettings.instance.json. Example:

```
{
  "Administration": {
    "SqlDbOptions": {
      "ConnectionString": "<connection string>",
      "LogSqlQueryParameterValues": true // Add this line to your
      ↪ config file to log SQL query param values for your SQL Server Administration database
    }
  },
  // OR:
  {
    "SqliteDbOptions": {
      "ConnectionString": "<connection string>",
      "LogSqlQueryParameterValues": true // Add this line to your

```

(continues on next page)

(continued from previous page)

```
↪config file to log SQL query param values for your Sqlite Data database
    }
}
}
```

## FIRELY SERVER DEPLOYMENT OPTIONS

Firely Server can be deployed as:

- Binaries, see *Getting Started*
- As Docker image, see *Using Firely Server on Docker*
- With a helm chart, see *Using Firely Server on Kubernetes with a Helm chart*
- With *Firely Server for your Simplifier project*

The Binaries can be deployed on your own machines and on cloud services. For deployment on Azure we included *instructions*.

In order to test the *Access control and SMART on FHIR* you will need to *Set up an Identity Provider*.

### 6.1 Using Firely Server on Docker

We have created a Docker image for Firely Server, so you can run the server in any environment that supports Docker. For this section we assume you have Docker installed on your system. If you want to install Docker for Windows, please read `docker_win` for specific installation details.

#### 6.1.1 Getting started

Before you can run Firely Server, you will need to pull the Docker Firely Server container and request a license.

1. Open your favourite command line tool and execute this command: `> docker pull firely/server`
2. Go to the [Simplifier website](#), login and download your evaluation license.
3. Create a working directory for Firely Server and place the license file there.

**Warning:** If you use Docker, you may want to run multiple instances of Firely Server (e.g. with Kubernetes). Read *Running imports with multiple instances* for caveats with the Administration endpoint.

---

**Note:** Since Firely Server version 4.7.0, the Docker image name has changed from `simplifier/vonk` to `firely/server`. The old image name will be maintained for a few months to allow for a smooth transition. When updating to version 4.7.0, you should start to use the new image name. Versions 4.6.2 and older will stay available (only) on ‘simplifier/vonk’.

---

### 6.1.2 Running a Docker Firely Server in SQLite mode

The easiest and the default way to run a Docker Firely Server container is to run Firely Server in SQLite repository mode. Note that this is not the most performant mode - see MongoDB and SQL Server options below for that.

Open your command prompt and execute this command: `> docker images firely/server`

You will get a list that looks like:

```
PS C:\> docker images firely/server
```

REPOSITORY	TAG	IMAGE ID	CREATED	SIZE
firely/server	latest	b36769e2fe2e	2 weeks ago	755MB

Navigate to your working directory for Firely Server and run the container with this command:

- in cmd.exe: `docker run -d -p 8080:4080 --name firely.server -v %CD%/firelyserver-license.json:/app/firelyserver-license.json firely/server`
- in Powershell: `docker run -d -p 8080:4080 --name firely.server -v ${PWD}/firelyserver-license.json:/app/firelyserver-license.json firely/server`

If your license file has a different name, use that name instead of `firelyserver-license` on the left side of the `-v` parameter in the command above. E.g. `-v ${PWD}/my-fancy-license.json:/app/firelyserver-license.json`

This will spin up a Firely Server container. It maps the host port 8080 to the container port 4080 with the switch `-p 8080:4080`. It will give the container the name `firely.server` with the switch `--name firely.server`. Furthermore it mounts the local licensefile into the `/app` directory, which is the directory where Firely Server resides in the container. Finally it will run the container in background mode with the switch `-d`.

To test whether the container is running correctly, type the command: `> docker ps`

```
PS C:\run\vonk> docker ps
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS	NAMES
aa709bb96599	simplifier/vonk	"dotnet Vonk.Serve..."	8 seconds ago	Up 7 seconds	0.0.0.0:8080->4080/tcp	vonk.server

You can also take a look at the logs for Firely Server with: `> docker logs firely.server`

Open a browser and use the address `http://localhost:8080/`. This will show the landing page of Firely Server.

To stop the container just type: `> docker stop firely.server` And to start it again: `> docker start firely.server` To completely remove the container: `> docker rm firely.server`

### Adjust settings when running in a Docker container

In the paragraph above we showed how to use `-v` to mount the license file into the `app` directory. We can use the same technique to mount a custom settings file or logsettings file. Or a directory to host the SQLite database outside of the container. Or have the logfile written to your working directory. Below is a full example using PowerShell. For creating `appsettings.instance.json` refer to [Firely Server settings](#). For creating `logsettings.instance.json` refer to [Log settings](#).

**appsettings.instance.json**

```
{
  "SQLiteDbOptions": {
    "ConnectionString": "Data Source=./resourcedata/vonkdata.db;Cache=Shared",
    "AutoUpdateDatabase": true
  }
}
```

**logsettings.instance.json**

Settings for console and levels not included for brevity.

```
{
  "Serilog": {
    "WriteTo": [
      {
        "Name": "Async",
        "Args": {
          "configure": [
            {
              "Name": "File",
              "Args": {
                "path": "./log/vonk.log",
                "rollingInterval": "Day",
                "fileSizeLimitBytes": "",
                "retainedFileCountLimit": "7",
                "outputTemplate": "{Timestamp:yyyy-MM-dd HH:mm:ss.fff zzz}
→{UserId} {Username} [{Application}] [{Level}] [Machine: {MachineName}] [ReqId:
→{RequestId}] {Message}{NewLine}{Exception}",
                "restrictedToMinimumLevel": "Information"
              }
            }
          ]
        }
      }
    ],
    "Enrich": [ "FromLogContext", "WithMachineName", "WithThreadId" ],
    "Properties": {
      "Application": "Firely Server",
      "Environment": "Default"
    }
  }
}
```

### Powershell

```
mkdir logs
mkdir resourcedata //do not use 'data' - the administration database is already in that
↳ folder in the container
//create the appsettings.instance.json above
//create the logsettings.instance.json above

docker run -d -p 8080:4080 --name firely.server `
-v ${PWD}/firelyserver-license.json:/app/firelyserver-license.json `
-v ${PWD}/appsettings.instance.json:/app/appsettings.instance.json `
-v ${PWD}/logsettings.instance.json:/app/logsettings.instance.json `
-v ${PWD}/resourcedata:/app/resourcedata `
-v ${PWD}/log:/app/log `
firely/server
```

You should see a vonkdata.db appear in the ./resourcedata folder, and a log file in the ./log folder. From here you can experiment with other settings. You can also easily keep different settings files side-by-side, mapping the one you want to test into the container, e.g. `-v ${PWD}/some-weird-settings.json:/app/appsettings.instance.json`.

### Mounting a your custom plugins folder into the container

In a similar way as described above, you can mount your custom plugins into the container. However, this will replace the default plugins; they need to be added manually.

- Download the binaries of your server version from <https://downloads.simplifier.net/firely-server/versions/>
- Create a plugin folder, including the default plugins (extracted from the downloaded archive)
- Create `appsettings.instance.json` with the appropriate Pipeline section (see *Configure the pipeline*)
- Mount the files and folder as shown below

```
docker run -d -p 8080:4080 --name firely.server `
-v ${PWD}/firelyserver-license.json:/app/firelyserver-license.json `
-v ${PWD}/appsettings.instance.json:/app/appsettings.instance.json `
-v ${PWD}/plugins:/app/plugins `
firely/server
```

The server is now accessible on `http://localhost:8080/`.

### Spinning up with a docker-compose file

Another way to spin up a Firely Server container is to use a docker-compose file. The above example can also be established by the following `docker-compose-sqlite.yml`:

```
1 version: '3'
2
3 services:
4
5     vonk-web:
6         image: firely/server
```

(continues on next page)

(continued from previous page)

```

7  ports:
8    - "8080:4080"
9  environment:
10   - VONK_Repository=SQLite
11   - VONK_Administration:Repository=SQLite
12   - VONK_License:LicenseFile=./license/firelyserver-license.json
13  volumes:
14   - ./app/license

```

Save the text above to a file in your working directory with the name `docker-compose.sqlite.yml` and then run the following command: `> docker-compose -f docker-compose.sqlite.yml up -d`

If your license file has a different name, use that name instead of `firelyserver-license` in the text above - but make sure to keep `./license` as that maps to a Docker volume inside the container.

```

PS C:\run\vonk> docker-compose -f docker-compose.memory.yml up -d
Creating vonk_vonk-web_1
PS C:\run\vonk> docker ps
CONTAINER ID   IMAGE          COMMAND                  CREATED        STATUS        PORTS                    NAMES
2c99a0648c3e   simplifier/vonk   "dotnet Vonk.Serve..."   5 seconds ago   Up 4 seconds   0.0.0.0:8080->4080/tcp    vonk_vonk-web_1
PS C:\run\vonk>

```

To stop the container, run: `> docker-compose -f docker-compose.sqlite.yml down`

**Note:** Strictly the settings for `VONK_Repository` and `VONK_Administration` are not needed here, since SQLite is the default setting. The settings are included to show where to configure the type of database to use. Much more information on that topic is in the paragraphs below.

**Note:** If you want to use a healthcheck in your docker-compose file, see [Liveness and readiness](#).

### 6.1.3 Running Docker with a SQL Server container

Firely Server can use also other repositories than Memory, for example SQL Server. This section describes how to spin up a Firely Server container and a SQL Server container. We will use docker-compose to achieve this.

**Warning:** SQL Server container uses at least 3.25 GB of RAM. Make sure to assign enough memory to the Docker VM if you're running on Docker for Mac or Windows.

**Warning:** If you also run the Administration database on SQL Server, please read [Retain the import history](#).

```

1  version: '3'
2
3  services:
4    vonk-web:
5      image: firely/server
6      ports:
7        - "8080:4080"
8      depends_on:

```

(continues on next page)

(continued from previous page)

```

9     - vonk-sqlserver-db
10    environment:
11      - VONK_Repository=SQL
12      - VONK_SqlDbOptions:ConnectionString=Initial Catalog=VonkData;Data Source=vonk-
↪sqlserver-db,1433;User ID=sa;Password=SQLServerStrong(!)Password*
13      - VONK_SqlDbOptions:SchemaName=vonk
14      - VONK_SqlDbOptions:AutoUpdateDatabase=true
15      - VONK_SqlDbOptions:AutoUpdateConnectionString=Initial Catalog=VonkData;Data_
↪Source=vonk-sqlserver-db,1433;User ID=sa;Password=SQLServerStrong(!)Password*
16      - VONK_Administration:Repository=SQL
17      - VONK_Administration:SqlDbOptions:ConnectionString=Initial Catalog=VonkAdmin;Data_
↪Source=vonk-sqlserver-db,1433;User ID=sa;Password=SQLServerStrong(!)Password*
18      - VONK_Administration:SqlDbOptions:SchemaName=vonkadmin
19      - VONK_Administration:SqlDbOptions:AutoUpdateDatabase=true
20      - VONK_Administration:SqlDbOptions:AutoUpdateConnectionString=Initial_
↪Catalog=VonkAdmin;Data Source=vonk-sqlserver-db,1433;User ID=sa;
↪Password=SQLServerStrong(!)Password*
21      - VONK_License:LicenseFile=./license/firelyserver-license.json
22    volumes:
23      - ./app/license
24
25    vonk-sqlserver-db:
26      image: mcr.microsoft.com/mssql/server
27      ports:
28        - "1433:1433"
29      environment:
30        - ACCEPT_EULA=Y
31        - SA_PASSWORD=SQLServerStrong(!)Password*
32      healthcheck:
33        test: /opt/mssql-tools/bin/sqlcmd -S localhost -U sa -P 'SQLServerStrong(!
↪)Password*' -Q 'SELECT 1 FROM VonkData.sys.tables'
34        interval: 1m30s
35        timeout: 10s
36        retries: 3

```

Save the text above to a file in your working directory with the name `docker-compose.mssqlserver.yml`. Make sure your Firely Server license file is named `firelyserver-license.json` and is residing in your working directory (see *Getting started* on how to obtain the license), **not** in a subdirectory named `license` (that is an internal directory inside the container). If your license file has a different name, use that name instead of `firelyserver-license` in the text above.

**Warning:** Without specifying `MSSQL_PID`, the server will run SQL Server Developer Edition. For production, you are responsible for obtaining the appropriate license (and setting the value accordingly). See [https://hub.docker.com/\\_/microsoft-mssql-server](https://hub.docker.com/_/microsoft-mssql-server) for more information.

**Note:** The current configuration will use SQL Server 2019. There are other versions of SQL server available, which can be selected by adding a tag to the value of `vonk-sqlserver-db:image` (e.g. `mcr.microsoft.com/mssql/server:2022-latest`) in the `yml`-file. We have successfully tested 2017-latest, 2019-latest and 2022-latest.

Then use this command to spin up a Firely Server container and SQL container: `> docker-compose -f`



```
docker-compose.mssqlserver.yml up -d
```

Open a browser and use the address <http://localhost:8080/>. This will show the landing page of Firely Server.

**Warning:** Wait for about 2 minutes, because it takes a while to fire up the SQL container

### 6.1.4 Running Docker with a SQL Server on host

Another possibility is to run a Firely Server container with a SQL Server repository on the host. You will need a Microsoft SQL Server running on your host. The version of SQL Server must at least be version 2012.

**Warning:** If you also run the Administration database on SQL Server, please read [Retain the import history](#).

To run the Firely Server container we will use the following docker-compose file:

```

1 version: '3'
2
3 services:
4
5   vonk-web:
6     image: firely/server
7     ports:
8       - "8080:4080"
9     environment:
10      - VONK_Repository=SQL
11      - VONK_SqlDbOptions:ConnectionString=Database=VonkData;Server=my_host\
    ↪ <myInstanceName>;User ID=<myUser>;Password=<myPassword>
12      - VONK_SqlDbOptions:SchemaName=vonk
13      - VONK_SqlDbOptions:AutoUpdateDatabase=true
14      - VONK_SqlDbOptions:AutoUpdateConnectionString=Database=VonkData;Server=my_host\
    ↪ <myInstanceName>;User ID=<DLLUser>;Password=<myPassword>
15      - VONK_Administration:Repository=SQL
16      - VONK_Administration:SqlDbOptions:ConnectionString=Database=VonkAdmin;Server=my_
    ↪ host\<myInstanceName>;User ID=<myUser>;Password=<myPassword>
17      - VONK_Administration:SqlDbOptions:SchemaName=vonkadmin
18      - VONK_Administration:SqlDbOptions:AutoUpdateDatabase=true
19      - VONK_Administration:SqlDbOptions:AutoUpdateConnectionString=Database=VonkAdmin;
    ↪ Server=my_host\<myInstanceName>;User ID=<DLLUser>;Password=<myPassword>
20      - VONK_License:LicenseFile=./license/firelyserver-license.json
21     volumes:
22      - ./app/license
23     extra_hosts:
24      - "my_host:192.0.2.1"

```

Save the text above to a file in your working directory with the name `docker-compose.mssqlserver_host.yml`. Before we spin up the container we have to adjust the `docker-compose.mssqlserver_host.yml`:

- On lines 11, 14, 16 and 19 the connection string to the database server is stated. Change the `Server` to your database server and instance name.
- Also change the `User ID` and `Password` on lines 11, 14, 16 and 19 to your credentials.

- Furthermore we have to tell Docker which IP address the host uses. This is done on line 24. In this case the host (named my\_host) uses IP address 192.0.2.1. Change this to the appropriate address.

After saving your settings, make sure your Firely Server license file is named `firelyserver-license.json` and is residing in your working directory (see [Getting started](#) on how to obtain the license), **not** in a subdirectory named `license` (that is an internal directory inside the container). Or use the name of your license file instead of `firelyserver-license` in the text above.

You can run the Firely Server container as follows: `> docker-compose -f docker-compose.mssqlserver_host.yml up -d`

A database will automatically be created if it is not already present on the database server. See [this page](#) for an overview of permissions the database user needs for creating the database and/or schema.

Open a browser and use the address <http://localhost:8080/>. This will show the landing page of Firely Server.

**Warning:** When you have a firewall installed on your host machine, it can block traffic from your Firely Server container to your host. Provide an inbound rule to allow traffic from the container to the host.

### 6.1.5 Run Docker with a MongoDB container

This section describes how to spin up a Firely Server container and a MongoDB container using a `docker-compose`. We assume you already have MongoDB installed.

**Warning:** If you also run the Administration database on MongoDB, please read [Retain the import history](#).

```
1 version: '3'
2
3 services:
4
5     vonk-web:
6         image: firely/server
7         environment:
8             - VONK_Repository=MongoDb
9             - VONK_MongoDbOptions:ConnectionString=mongodb://vonk-mongo-db/vonkdata
10            - VONK_MongoDbOptions:EntryCollection=vonkentries
11            - VONK_Administration:Repository=MongoDb
12            - VONK_Administration:MongoDbOptions:ConnectionString=mongodb://vonk-mongo-db/
13      ↪vonkadmin
14            - VONK_Administration:MongoDbOptions:EntryCollection=vonkentries
15            - VONK_License:LicenseFile=./license/firelyserver-license.json
16         volumes:
17             - ./app/license
18         ports:
19             - "8080:4080"
20
21     vonk-mongo-db:
22         image: mongo
```

Save the text above to a file in your working directory with the name `docker-compose.mongodb.yml`. Make sure your Firely Server license file is named `firelyserver-license.json` and is residing in your working directory (see [Getting started](#) on how to obtain the license), **not** in a subdirectory named `license` (that is an internal directory inside

the container). If your license file has a different name, use that name instead of `firelyserver-license` in the text above.

Use this command to spin up a Firely Server container and MongoDB container: `> docker-compose -f docker-compose.mongodb.yml up -d`

Open a browser and use the address <http://localhost:8080/>. This will show the landing page of Firely Server.

### 6.1.6 Providing license via an environment variable

The examples above demonstrate how to provide a license by mounting a license file from the host's filesystem to the filesystem within a container. Starting from Firely Server v4.7.0, the license can also be provided as a string via an environment variable named `VONK_License:LicenseString`. This is meant to simplify deployments of Firely Server within Docker as you don't have to store the license file on the host's filesystem anymore.

The value of that variable should contain the same text as the license file but all in one line.

The example below shows how to spin up a Docker container by supplying the license as a variable.

```
docker run -d \
  -p 8080:4080 \
  --name firely.server \
  -e "VONK_License:LicenseString={ 'LicenseOptions': { 'Kind': 'Production', 'ValidUntil'
  ↪ ': '2022-10-30', 'Licensee': 'example@fire.ly', 'Plugins': [ ... ] }, 'Signature': '...
  ↪ ' }" \
  firely/server
```

If you use docker-compose, you can specify the variable in you docker-compose file like this:

```
1 version: '3'
2
3 services:
4
5   vonk-web:
6     image: firely/server
7     ports:
8       - "8080:4080"
9     environment:
10      - "VONK_License:LicenseString={ 'LicenseOptions': { 'Kind': 'Production',
  ↪ 'ValidUntil': '2022-10-30', 'Licensee': 'example@fire.ly', 'Plugins': [ ... ] },
  ↪ 'Signature': '...' }"
```

## 6.2 Using Firely Server on Kubernetes with a Helm chart

It is very common to run Firely Server on a Kubernetes cluster. This can be any Kubernetes implementation, it is not bound to a specific cloud provider.

To ease the process we provide a helm chart to deploy Firely Server on the cluster. The chart can be found in the [FirelyTeam/Helm.Charts](#) repository.

All the settings are described in the [Read.me](#).

The same repository also contains a helm chart to deploy *Identity Server*, but that is very early still. Documentation and settings will be added to that as well.

### 6.3 Firely Server deployment on Azure Web App Service

In this section we explain how you can deploy Firely Server in the Azure cloud.

#### 6.3.1 Getting started

Before you can run Firely Server, you will need to download the Firely Server binaries and request a license:

1. Go to the Simplifier website, login and download the Firely Server binaries from <https://simplifier.net/vonk/download>
2. Download the trial license file from the same location.

#### 6.3.2 Deployment

1. Go to Azure (<https://portal.azure.com>) and create a web app:

Create

- ## 6.3. Firely Server deployment on Azure Web App Service 135

create a new one and select Windows for the operation system (OS):

## Create Web App

**Basics**   Monitoring   Tags   Review + create

App Service Web Apps lets you quickly build, deploy, and scale enterprise-grade web, mobile, and API apps running on any platform. Meet rigorous performance, scalability, security and compliance requirements while using a fully managed platform to perform infrastructure maintenance. [Learn more](#)

### Project Details

Select a subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription \* ⓘ Visual Studio Professional Subscription ▼

Resource Group \* ⓘ (New) firelyserver ▼

[Create new](#)

### Instance Details

Name \* firelyserver ✓ .azurewebsites.net

Publish \* ☒ Code ☐ Docker Container

Runtime stack \* .NET Core 3.1 (LTS) ▼

Operating System \* ☐ Linux ☒ Windows

3. Add the trial license file (firelyserver-trial-license.json) to the firely-server-latest.zip by dragging the license file into the zipfile.
4. Open a webbrowser, navigate to <https://<webapp>.scm.azurewebsites.net/ZipDeployUI> and drag vonk\_distribution.zip into the browser window. This will install the Firely Server as a Web App in Azure. In our example the url is <https://firelyserver.scm.azurewebsites.net/ZipDeployUI> This method of deployment does not work in Internet Explorer. It does work in Firefox, Chrome and Edge. Please make sure that after you have uploaded the .zip file, all content is extracted into the top-level webroot directory.


/wwwroot + | 184 items

	sqlserver	1/2/2022, 7:18:24 PM	
	vonk-import	1/2/2022, 7:19:09 PM	
	vonk-import.R4	1/2/2022, 7:19:17 PM	
	appsettings.default.json	1/2/2022, 7:25:02 PM	14 KB
	Dapper.dll	1/2/2022, 7:24:54 PM	193 KB
	DnsClient.dll	1/2/2022, 7:24:54 PM	136 KB
	errataFhir3.0.zip	1/2/2022, 7:25:04 PM	16 KB
	errataFhir4.0.zip	1/2/2022, 7:25:04 PM	4 KB

5. Open a browser and go to the site <https://<webapp>.azurewebsites.net/>. This will show the Firely Server home page.

### 6.3.3 Change database

In this example Firely Server is using a memory repository. If you want to change it to another kind of repository then you could change that on the page Application Settings of the Web App. Here you can set *Environment Variables* with the settings for either *SQL Server* or *MongoDB*. For example for MongoDB it will look like this:



### 6.3.4 More information

About Azure zip deployment: <https://docs.microsoft.com/en-us/azure/app-service/app-service-deploy-zip#deploy-zip-file>

## 6.4 Firely Server for your Simplifier project

When you are interested in using a project on Simplifier for you conformance resources, you can now quickly spin up a Firely Server instance with the content of the project already loaded. For more information how to do this please see [Use Firely Server with your Simplifier artifacts..](#)

## 6.5 Deploy Firely Server on a reverse proxy

### 6.5.1 Why

For ASP.NET Core 1.0 Microsoft suggested to always use another web server in front of Kestrel for public websites. For ASP.NET Core 2.0, while this is not a hard constraint anymore there are still a series of advantages in doing so:

- some scenarios like sharing the same IP and port by multiple applications are not yet supported in Kestrel
- helps in limiting the exposed surface area
- provides an additional layer of configuration and defense
- provides process management for the ASP.NET Core application (ensuring it restarts after it crashes)
- in some scenarios a certain web server already integrates very well
- helps simplifying load balancing and SSL setup

Hence using a reverse proxy together with the Kestrel server allows us to get benefits from both technologies at once.

### 6.5.2 With IIS

A common option on Windows is using IIS: for instructions on how to deploy Firely Server on IIS see *Deploy Firely Server on IIS*.

For a comparison of IIS and Kestrel features at the moment of this writing you can [check here](#).

### 6.5.3 With Nginx

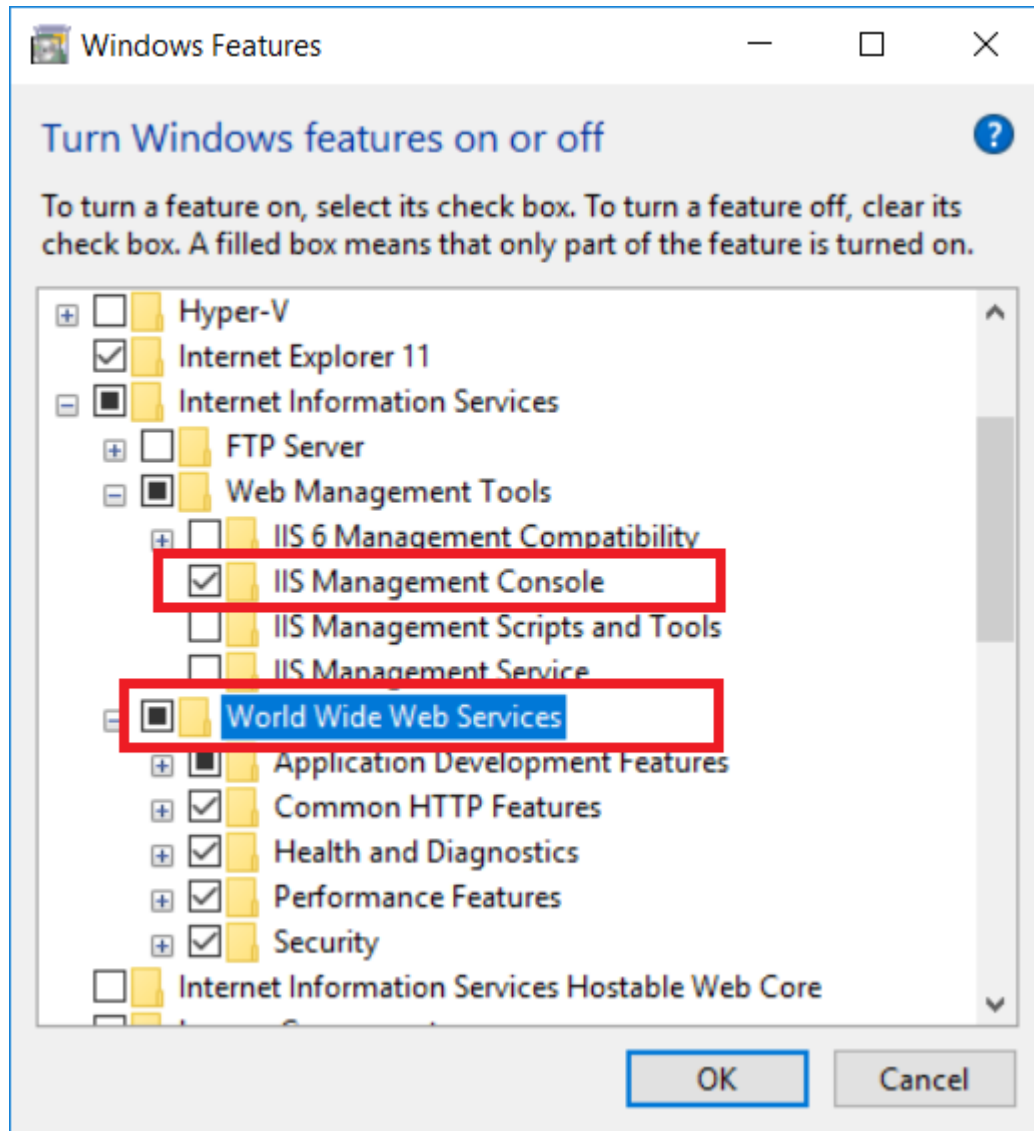
A popular open source alternative is Nginx. For instruction on how to deploy Firely Server on Nginx see *Deploy Firely Server on Nginx on Linux*.

## Deploy Firely Server on IIS

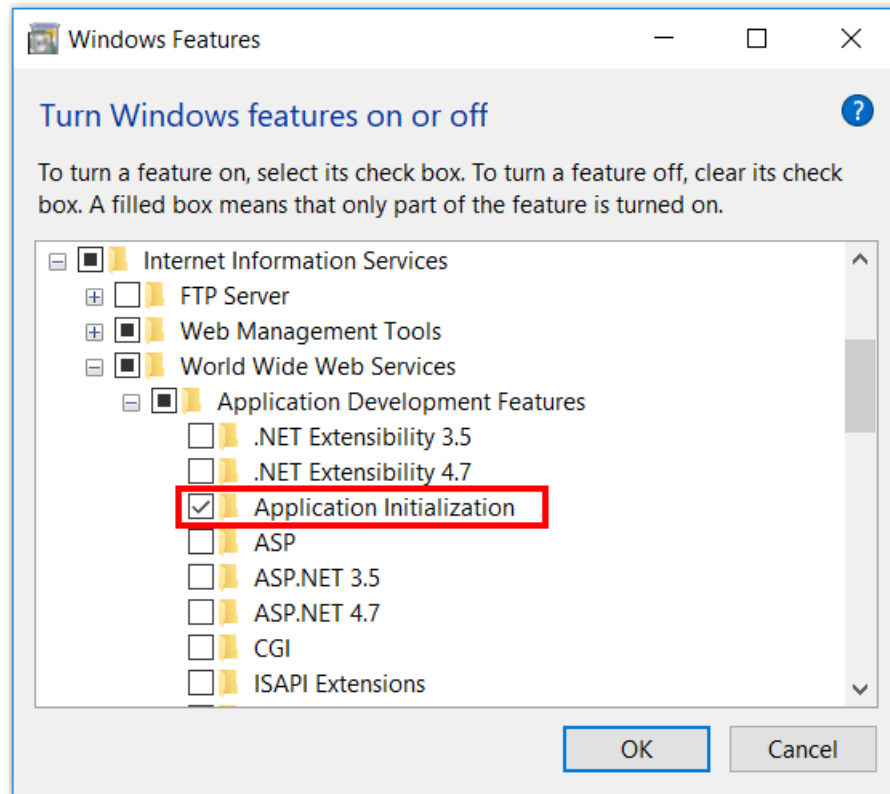
### Prerequisites

1. The following operating systems are supported: Windows 7 or Windows Server 2008 R2 and later
2. Have IIS windows feature turned on on the hosting machine. You can do this in **Control Panel -> Programs and Features -> Turn Windows features on or off** . You need to check **Internet Information Services -> Web Management Tools -> IIS Management Console** and **Internet Information Services -> World Wide Web Services** to accept the default features or customize IIS to fit your needs.






3. In order to have IIS start the Firely Server right away, you will have to enable the “Application Initialization” setting:



4. To enable PUT and DELETE interactions, you will have to turn off WebDAV:



See <https://stackoverflow.com/questions/6739124/iis-7-5-enable-put-and-delete-for-restful-service-extensionless> for some background information. If you do not want to disable WebDAV for all of IIS, you can also disable it just for Firely Server using a setting in *web.config*, see *Configuration*.

5. Choose a solution to deploy/move the application to the hosting system. Multiple alternatives exist like Web Deploy, Xcopy, Robocopy or Powershell. One popular choice is using Web Deploy in Visual Studio. For using that you will need to install Web Deploy on the hosting system. To install Web Deploy, you can use the Web Platform Installer (<https://www.microsoft.com/web/downloads/platform.aspx>).
6. Install the *.NET Core Runtime & Hosting bundle* on the hosting system. After installing it, you may need to do a “net stop was /y” and “net start w3svc” to ensure all the changes are picked up for IIS. The bundle installs the .NET Core Runtime, .NET Core Library, and the ASP.NET Core Module. ASP.NET Core Module (ANCM) allows you to run ASP.NET Core applications using Kestrel behind IIS. For more information about ANCM check <https://docs.microsoft.com/en-us/aspnet/core/fundamentals/servers/aspnet-core-module>

6. Prepare binaries. You can either download the binaries for Firely Server (see *Getting Started*), or create your own solution by building a facade.

If you are deploying a Firely Server facade in your own web server application, take the additional following prerequisites into consideration:

- Make sure you use the **IISIntegration NuGet** package. You can install this as part of one of the metapackages (`Microsoft.AspNetCore` and `Microsoft.AspNetCore.All`) or independently `Microsoft.AspNetCore.Server.IISIntegration`. This is needed for the interoperability between Kestrel and ANCM.
- Provide a *web.config* file for configuring the ANCM module or make sure the selected deploy options generates one for you. Using **dotnet publish** or **Visual studio publish** would generate a *web.config* for you. Check <https://docs.microsoft.com/en-us/aspnet/core/hosting/aspnet-core-module> for guidance on configuring the ANCM.

### Create Website in IIS

1. Publish the application to a location on the host server, using the solution selected in the Prerequisites step 3.
2. In IIS Manager create a new website or a new application under existing IIS site. Fill the **Site name**, the **Binding** and link the **Physical path** to the folder created in the previous step, similar to the picture below. The bindings defined in IIS override any bindings set in the application by calling either `Listen` or `UseUrls`.

Add Website

Site name: Vonk\_IIS Application pool: Vonk\_IIS Select...

Content Directory

Physical path: D:\Project\Vonk ...

Pass-through authentication

Connect as... Test Settings...

Binding

Type: http IP address: All Unassigned Port: 80


Host name: vonk.mydomain.org

Example: www.contoso.com or marketing.contoso.com

☐ Start Website immediately

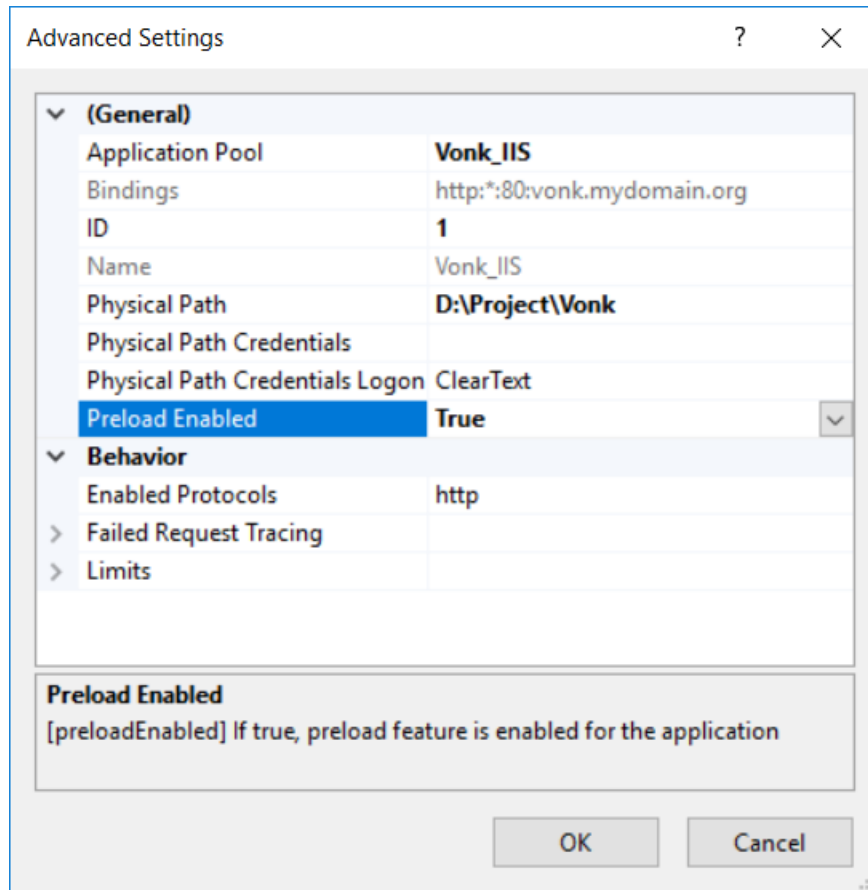
OK Cancel

3. Edit the application pool to set the **.NET CLR VERSION** to **NO Managed Code**, similar to the picture below (we use IIS as a reverse proxy, so it isn't actually executing any .NET code). To edit the application pool, go to the **Application Pools** panel, right-click the website's app pool and select **Advanced Settings...** from the popup menu.



**Attention:** IIS will pause a website if it is idle for a while. Pausing a dotnet process is the same as shutting it down, so this means that IIS shuts down the Firely Server. This causes problems with each first request that is sent to Firely Server after an idle period. Firely Server needs a couple of seconds to start up again, and answers with a 423 Lock Error while it loads. Make sure to set **Start Mode** to **AlwaysRunning** to prevent IIS from shutting down Firely Server.

4. Go to your IIS Sites, and click on your Firely Server website. On the right, choose **Advanced settings...**. Set **Preload Enabled** to **True** to make IIS load the Firely Server right away, instead of on the very first request. Otherwise this first request results in a 423 Lock Error as described above.



## Configuration

- You can use *web.config* to configure ASP.NET Core Module and IIS using the `<system.webServer>` section. Read more about configuring ANCM at <https://docs.microsoft.com/en-us/aspnet/core/hosting/aspnet-core-module>.
- If you need to disable WebDAV for Firely Server, so you can perform PUT and DELETE interactions, add this to the *web.config* file:

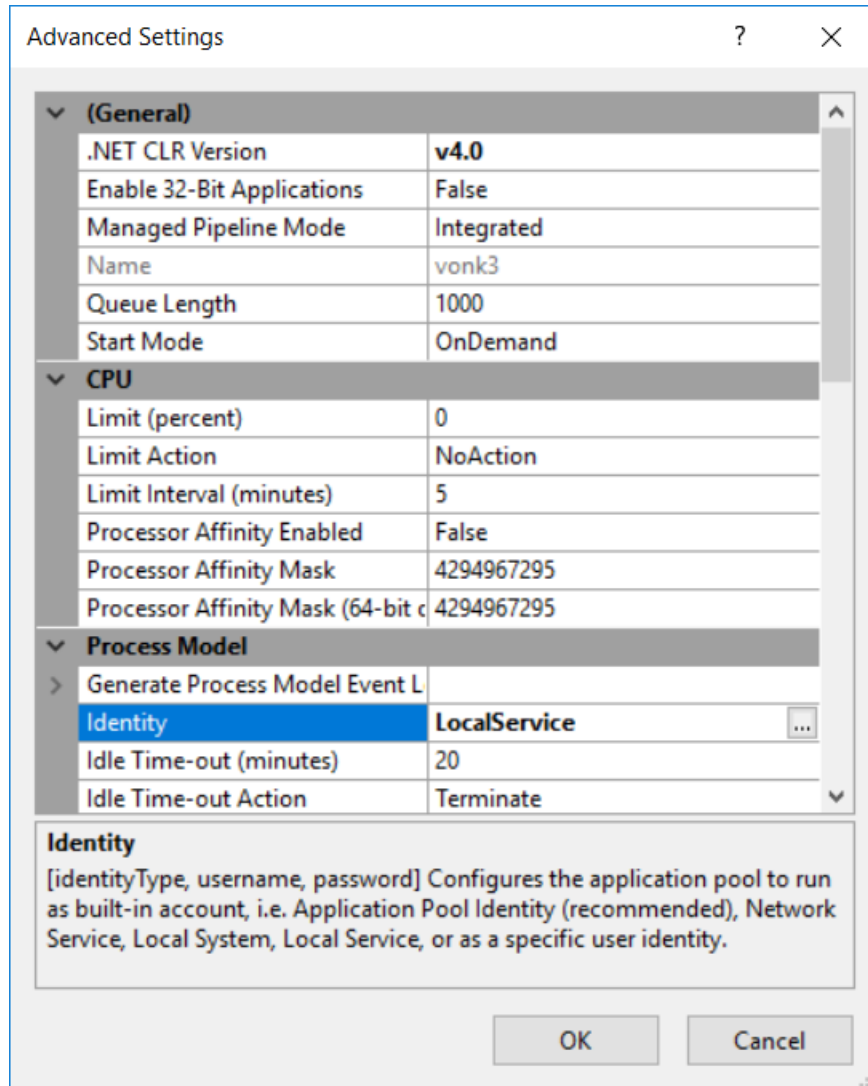
```
<modules>
  <remove name="WebDAVModule" />
</modules>
<handlers>
  <remove name="WebDAV" />
</handlers>
```

- You can configure the Firely Server using the *appsettings.json* file (see *Configuring Firely Server*).
- If you are building a Firely Server facade, you can use *IISOptions* to configure *IISIntegration* service options. You can read more about *IISOptions* at <https://docs.microsoft.com/en-us/dotnet/api/microsoft.aspnetcore.builder.iisoptions?view=aspnetcore-2.0>.

```
services.Configure<IISOptions>(options =>
{
    ...
});
```

## SQL

In order to use the Sql Repository option in IIS you should make sure that the identity of the IIS application pool has rights to use the database considering the provided connection string. To change the identity the application pool is using open IIS Application Pools select your application pool right click and select “Advanced Settings...” You should see something similar to the image below:



## Extra

If you'd like to set Firely Server environment variables via an Azure pipelines task, you can do so by setting the application pool's environment variables. For example, to pass the variable TEST to Firely Server that's housed in the fhir application pool, do the following:

```
%systemroot%\system32\inetsrv\APPCMD set config -section:system.applicationHost/
↪ applicationPools /- "[name='fhir'].environmentVariables.[name='VONK_TEST']" /
↪ commit:apphost
%systemroot%\system32\inetsrv\APPCMD set config -section:system.applicationHost/
```

(continues on next page)



(continued from previous page)

```
↪applicationPools /+"[name='fhir'].environmentVariables.[name='VONK_TEST',value='some_
↪value_here']" /commit:apphost
```

See also *Firely Server settings with Environment Variables*.

## Deploy Firely Server on Nginx on Linux

### About Nginx

NGINX is a popular open source web server. It can act as a reverse proxy server for TCP, UDP, HTTP, HTTPS, SMTP, POP3, and IMAP protocols, as well as a load balancer and a HTTP cache. You can find the documentation for the Nginx server at <https://nginx.org/en/docs/>.

### Prerequisites

1. The following linux distribution are supported: Ubuntu, RHEL, Debian, Fedora, CentOS, SLES
2. Install .Net Core on the machine (see <https://www.microsoft.com/net/learn/get-started/linuxubuntu>)
3. Install Nginx `sudo apt-get install nginx`

### Start Kestrel Firely Server

Download the binaries for Firely Server (see *Getting Started*), open a terminal console and start the Firely Server process by using: `dotnet Vonk.Server.dll`. You should be able to reach to home page at `http://localhost:4080` (or a different port if you changed the default configurations)

### Configure Nginx as a reverse proxy

To configure Nginx as a reverse proxy to forward requests to our ASP.NET Core application, modify `/etc/nginx/sites-available/default`. Open it in a text editor, and replace the contents with the following:

```
server {
    listen 80;
    # Match incoming requests with the following path and forward them to
    # the location of the Kestrel server.
    # See http://nginx.org/en/docs/http/nginx_http_core_module.html#location
    location / {
        #This should match the location where you deployed the Firely Server binaries.
        ↪with the Kestrel server.
        #This can be on the same machine as the Nginx server or on a separate dedicated.
        ↪machine
        proxy_pass http://localhost:4080;
        # The Kestrel web server we are forwarding requests to only speaks HTTP 1.1.
        proxy_http_version 1.1;
        proxy_set_header Upgrade $http_upgrade;
        # Adds the 'Connection: keep-alive' HTTP header.
        proxy_set_header Connection keep-alive;
        # Forwards the Host HTTP header.
```

(continues on next page)

(continued from previous page)

```
    proxy_set_header Host $host;
    proxy_cache_bypass $http_upgrade;
}
}
```

Now you can run the Firely Server.

### Configuration

- To configure the Firely Server, you can use the appsettings.json file (see [Configuring Firely Server](#)).
- To configure Nginx you need to add extra options to the /etc/nginx/sites-available/default or to the nginx.conf file.
- To monitor the application you can use systemd and create a service for starting, stopping and managing the process.

#### 6.5.4 Using X-Forwarded-Host header

The [X-Forwarded-Host header](#) is used for identifying the original host by the client, which is especially useful with reverse proxies. As per the [FHIR specification](#), Firely Server supports usage of this header. Keep in mind that usage of the header exposes privacy sensitive information.

When using this header, make sure that the header value only contains the domain name like listed below:

##### Works:

- fire
- www.fire.ly
- www.fire

##### Does not work:

- fire/
- https://fire.ly

## 6.6 Set up an Identity Provider

### 6.6.1 About Identity Providers and Firely Server

In order to use [Access control and SMART on FHIR](#) you need an Identity Provider that can provide OAuth2 JWT Tokens with claims that conform to [SMART on FHIR](#). In a production scenario, you typically already have such a provider. It could be the EHR system, the Active Directory, or a provider set up specifically for let's say a Patient Portal. It is also very well possible that the provider handing the correct claims uses a federated OAuth2 provider to do the authentication.

## 6.6.2 Firely Auth

In order to provide a turn-key experience, Firely offers Firely Auth as an add-on. It provides an external authorization services optimized for SMART on FHIR. For more details, see [Firely Auth](#).

To allow you to test *Access control and SMART on FHIR*, we provide you with [instructions](#) to build and run Firely Auth in which you can configure the necessary clients, claims and users yourself to test different scenarios.

## 6.7 Performance of Firely Server

This section contains a general approach to Firely Server performance and SQL Server guidelines:

### 6.7.1 Performance of Firely Server

#### About Performance

What is the performance of Firely Server? That is a simple question, but unfortunately it has no answer. Or more precisely, the answer depends on many variables. On this page we try to give you insight into those variables, introduce you to testing performance yourself, and finally present the results of the tests that we run ourselves.

#### Performance variables

##### Firely Server Configuration

Firely Server can be run as self contained FHIR Server or as a Facade on top of an existing system. The performance of a Facade is heavily dependent on the performance of the system it is built on top of. The self contained server can run on different databases. On top of that you can configure a couple of features in the settings that influence the performance. These are the most important configuration variables to take into account:

#### Repository

1. Memory: Memory is only meant for quick tests, and for use in unittests. Do not use it in any serious scenario, much less for performance critical scenarios.
2. SQLite: SQLite is mainly used for the Administration database of Firely Server, but you can also use it for the main database. Deployment is very easy because of the zero footprint of the driver, but be aware of its limits. Firely Server must have very fast access to the database file, so effectively it has to be on a local disk. Multithreading does work, but does not scale as well as other databases.
3. SQL Server: Performance tuning of SQL Server is a topic on its own. Firely Server manages the tables it needs, and the indexes on top of it are optimized for the way Firely Server queries them. See [SQL Server Guidelines](#) for more details about SQL Server database maintenance and tuning.
4. MongoDB: Performance tuning of MongoDB is, as well, a topic on its own. Firely Server manages the collections it needs, and the indexes on top of it are optimized for the way Firely Server queries them. MongoDB is used in our own performance tests, see below.

### Prevalidation

By default, a resource that is sent to Firely Server is fully validated against its `StructureDefinition`. This requires extra processing and thus extra time. But you can disable full validation if needed, with the [ValidateIncomingResources](#) setting. We have no tests in place yet to time the difference caused by this setting.

### Search Parameters

When a resource is sent to Firely Server for storage, Firely Server indexes the resource for all the search parameters that are applicable to the resource. The more search parameters are known to Firely Server, the more resources will be used for indexing - both time (for the indexing processing) and storage (for storing the values of each search parameter). This also increases the size of the index tables and indexes, and therefore querying times. Thus, if you know you will only use a portion of the predefined search parameters you can choose to delete the others from the Administration API - see [Managing Conformance Resources](#) and [Restrict supported resources and SearchParameters](#).

Additionally, these search parameters can have an impact on the search performance:

1. `_total`: Searchbundles contain the element `total` which indicates the total number of resources that match the query's search parameters. Setting this parameter to `_total = none` results in faster searches as it saves the query that would generate the result of the aforementioned element.

### Pipeline

Firely Server is made up of a pipeline of plugins. You can leave out any plugin that you don't need - so if you don't need conditional processing (create, update, delete), just exclude them from the pipeline. Excluded plugins are not loaded and thus never executed - see [Configuring the Firely Server Pipeline](#).

### Platform

Firely Server can run on Windows, Linux and MacOS. Directly or in a Docker container. On real hardware, virtual machine, app service or Kubernetes cluster. And then you can choose the dimensions of the platform, scaling up (bigger 'hard'ware) or scaling out (more (virtual) machines) as you see fit. Each of these choices will influence performance.

Besides the way Firely Server is deployed, the way the database is deployed is an important factor. Firely Server needs a very low latency connection between the Firely Server process(es) and the database server. If you have configured [Application Insights](#), the calls to the database are recorded as separate dependencies so you can check whether this may be a bottleneck.

Firely Server is optimized for multithreaded processing. This means that it will fully benefit from extra processing cores, either in the same machine (multi core processor) or by adding additional machines (and thus processors).

Firely Server is fully stateless, so no data about connections is saved between requests. First of all this helps in scaling out, since you don't need thread affinity. On top of that this reduces the memory requirements of Firely Server.

## Usage patterns

How will Firely Server be used in your environment?

1. Mostly for querying, or rather for creating and updating resources? Altering resources requires more processing than reading them. Also see the comment on indexing and search parameters above.
2. How is the distribution of values in the resources that you query on? E.g. if you use only a few types of resources, query them just by tag and the resources have only about 5 different tags, calculating the number of results will take a lot of time. Using more finegrained distributed values to query on solves this.
3. With many individual resources or with (large) batches or transactions? Transactions take a lot longer to process and require more memory, proportionally to the number of resources in them. If many transactions are run in parallel, requests may queue up.
4. Many users with a low request rate each, or a few heavy users? Since Firely Server is stateless, this has little influence. The total request rate is what counts.

## Testing performance yourself

Because of all the variables mentioned above the best way to find out whether Firely Server's performance is sufficient for your use is: test it yourself.

We provide an evaluation license that you can use for any testing, including performance testing. See [Getting Started](#).

## Variables

Before you start testing, study the variables above and provide answers to them. Then you can configure your platform and your tests in a way that comes closest to the expected real use.

## Requests

You need a set of requests that you want to test. Based on your use case, identify the 5 (or more) most frequent requests. For extra realism you should provide the parameters to the requests from a dataset (like a .csv file with search parameter values).

## What to measure?

There are essentially two questions that you can investigate:

1. Given this deployment, (mix of) requests and an expected request rate, what are the response times?
2. Given this deployment and (a mix of) requests, how many requests can Firely Server handle before it starts returning time-outs?

Besides response times more insight can be gained by measuring the load on the server (processor / memory usage, disk and network latency, for both the Firely Server and the database server) as well as the machine you are generating the requests from (to ensure that is not bottlenecked).

Always make sure to use at least 2 separate machines for testing: one for Firely Server, and a separate one for generating the requests. Testing Firely Server on the same machine as you're generating requests from will make Firely Server compete with the load testing tool for resources which'll hamper the legitimacy of the test results.

Based on the answers you can retry with different parameters (e.g. add/remove hardware) to get a sense of the requirements for real use deployment.

### Data

Performance testing is best done with data as realistic to your situation as possible. So if you happen to have historic data that you can transform to FHIR resources, that is the best data to test with.

But if you don't have data of your own, you can use synthesized data. We use data from the Synthea project for our own tests. And we provide *Firely Server Ingest (FSI)* to upload the collection bundles from Synthea to Firely Server (or any FHIR Server for that matter).

If you build a Facade, the historical data is probably already in a test environment of the system you build the Facade on. That is a perfect start.

### Test framework

To run performance tests you need a framework to send the requests in parallel and measure the response times. Test automation is a profession in itself so we cannot go into much detail here. You can search for 'REST Performance test tools' to get some options.

### Available performance figures

We are in the process of setting up performance tests as part of our Continuous Integration and Deployment. Here we describe how this test is currently set up. Because of the beta phase this is in, the output is not yet complete nor fully reliable. Nevertheless we share the preliminary results to give you a first insight.

### Firely Server performance test setup

1. Configuration
  1. Repository: MongoDB, both for Administration and for the main database.
  2. Prevalidation: off
  3. Search parameters: support all types of resources and all search parameters from the FHIR specification.
  4. Pipeline: load all available plugins except authorization.
2. Platform
  1. Azure Kubernetes Service, 2 nodes.
  2. Each node: Standard F2s (2 vcpus, 4 GB memory), running Linux
  3. 1 MongoDB pod and 2 Firely Server pods, plus the Kubernetes manager
3. Usage pattern - we created a simple mix of requests
  1. Upload the first 100 Synthea bundles from the precalculated set, each collection bundle transformed to a Batch.
  2. A 'general' test, consisting of:
    1. Query Patient by name: GET {url}/Patient?name=...
    2. Query Patient by name and maximum age: GET {url}/Patient?name={name}&birthdate=ge{year}
    3. Query all Conditions: GET {url}/Condition

4. Query a Patient by identifier, with Observations: GET {url}/Patient?identifier={some identifier}&\_revinclude=Observation:subject
5. Query a Patient by identifier, with Observations and DiagnosticReports: GET {url}/Patient?identifier={some identifier}&\_revinclude=Observation:subject&\_revinclude=DiagnosticReport:p
3. Page through all the CarePlan resources: GET {url}/CarePlan?\_count=10, and follow next links.
4. Page through 1/5 of the Patient resources and delete them: DELETE {url}/Patient/{id}
5. 20 concurrent users, randomly waiting up to 1 second before issuing the next request.
6. Test run of 5 minutes
4. Test framework
  1. Locust for defining and running tests
  2. Telegraf agents for collection metrics
  3. InfluxDB for storing results
  4. Grafana for displaying results

## Test results

1. Upload: not properly timed yet.
2. General test: 75 percentile of response times around 200 ms. Note that the responses on queries with ‘\_revinclude’ contain over 30 resources on average, sometimes over 100.
3. Page through all CarePlan resources: 75 percentile of response times around 110 ms.
4. Delete patients: This test always runs with 40 concurrent users, and 75 percentile of response times are around 350ms. Note that in Firely Server a delete is essentially an update, since all old versions are retained.

## 6.7.2 SQL Server Guidelines

Performance tuning of SQL Server is a topic on its own. Firely Server manages the tables it needs, and the indexes on top of it are optimized for the way Firely Server queries them. In this section we offer general guidelines on how to configure and maintain your SQL database. This topic only covers the most common options and should by no means be regarded as exhaustive.


## Memory Settings

The total amount of memory needed by SQL Server depends on your use case and (expected) workload. SQL Server will claim memory from the OS when needed but will not give it back to the OS once claimed. If no limit is specified you will run the risk of starving the host server and OS of enough free memory to operate efficiently. It is therefore strongly advised to limit the maximum amount SQL Server is allowed to use.

To set a fixed amount of memory using SSMS<sup>1</sup>:

1. In Object Explorer, right-click a server and select Properties.
2. Click the Memory node.
3. Under Server Memory Options, enter the amount that you want for ‘Minimum server memory (in MB)’ and ‘Maximum server memory (in MB)’.

<sup>1</sup> [Server memory configuration options](#)



Use the default settings to allow SQL Server to change its memory requirements dynamically based on available system resources. It is recommended to set a max server memory as detailed above.

Windows Server with Desktop Experience requires a minimum memory size of 2 GB<sup>2</sup>, but 4 GB or more is recommended (more is better). Example: If your server has 64GB of memory, set the amount of memory in SQL Server to 60,000MB (i.e. 64GB - 4GB = 60GB)

## Disk Configuration

When optimizing SQL Server it is highly recommended to set the tempDB, the data disk and the log disk on separate drives (preferably SSDs)<sup>1</sup>.

## TempDB

Place the tempDB database on a separate disk and adjust the number of secondary data files according to how many (logical) processors the host features<sup>2</sup>. If the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files. For TempDB it usually is sufficient to create 8 equally sized data files. As a general rule of thumb:

- **Number of logical processors < 8**  
→ Set the number of data files equal to the number of logical processors
- **Number of logical processors >= 8**  
→ Set the number of data files to 8

<sup>2</sup> Memory requirements for Windows Server

<sup>1</sup> Place Data and Log Files on Separate Drives

<sup>2</sup> TempDB data files



## Data & Log files

Placing both data AND log files on the same device can cause contention for that device, resulting in poor performance. Placing the files on separate drives allows the I/O activity to occur at the same time for both the data and log files.

Also format your data disk to use 64-KB allocation unit size for all data files placed on a drive<sup>3</sup>. Use SSDs with high IOPS where possible.

## Backup Maintenance Plan

When Firely Server is used in conjunction with SQL Server, a maintenance plan should be in place. Furthermore, if the recovery log is not set to 'Simple', regular back-ups of the transaction logs should be made.

Microsoft advises that if a database uses either the full or bulk-logged [recovery model](#), you must back up the transaction log regularly enough to protect your data, and to prevent the [transaction log from filling](#). This truncates the log and supports restoring the database to a specific point in time.

For more information on how to back up the Transaction Log consult the official Microsoft documentation, which can be found [here](#).

## Index Maintenance

You can check the index fragmentation with the following script:

```
SELECT OBJECT_SCHEMA_NAME(i.object_id) AS schema_name,
       OBJECT_NAME(i.object_id) AS object_name,
       i.name AS index_name,
       i.type_desc AS index_type,
       100.0 * (ISNULL(SUM(rgs.deleted_rows), 0)) / NULLIF(SUM(rgs.total_rows), 0) AS avg_
↪fragmentation_in_percent
FROM sys.indexes AS i
INNER JOIN sys.dm_db_column_store_row_group_physical_stats AS rgs
ON i.object_id = rgs.object_id
   AND
   i.index_id = rgs.index_id
WHERE rgs.state_desc = 'COMPRESSED'
GROUP BY i.object_id, i.index_id, i.name, i.type_desc
ORDER BY schema_name, object_name, index_name, index_type;
```

Some guidelines<sup>1</sup>:

- **Do not assume that index maintenance will always noticeably improve your workload.**
- **If you observe that rebuilding indexes improves performance, try replacing it with updating statistics.** This may result in a similar improvement. In that case, you may not need to rebuild indexes as frequently, or at all, and instead can perform periodic statistics updates.
- Monitor index fragmentation and page density over time to see if there is a correlation between these values trending up or down, and query performance. If higher fragmentation or lower page density degrade performance unacceptably, reorganize or rebuild indexes. **It is often sufficient to only reorganize or rebuild specific indexes used by queries with degraded performance.** This avoids a higher resource cost of maintaining every index in the database.

<sup>3</sup> Best practice storage

<sup>1</sup> Index maintenance

- Establishing a correlation between fragmentation/page density and performance also lets you determine the frequency of index maintenance. **Do not assume that maintenance must be performed on a fixed schedule.** A better strategy is to monitor fragmentation and page density, and run index maintenance as needed before performance degrades to an unacceptable level.
- If you have determined that index maintenance is needed and its resource cost is acceptable, perform maintenance during low resource usage times, if any, keeping in mind that resource usage patterns may change over time.

There is also a general rule of thumb with regards to fragmentation (YMMV):

- < 10%: no action needed.
- 10-30%: reorganize the index
- >30%: rebuild the index

### SQL Server Statistics

The maintenance of statistics falls outside the scope of Firely Server.

Normally SQL Server will update and maintain the statistics automatically. But when the database becomes large this process might happen too infrequently. There is no magic bullet for updating statistics, so you are advised to experiment and test what works best for your setup.

Consider updating statistics for the following conditions<sup>1</sup>:

- Query execution times are slow.
- Insert operations occur on ascending or descending key columns.
- After maintenance operations.

To manually update statistics there are two options<sup>2,3</sup>:

- A stored procedure `sp_updatestats`: this runs `UPDATE STATISTICS` against all user-defined and internal tables in the current database.
- A SQL statement against a single table (here we use the Entry table) `UPDATE STATISTICS [vonk].[entry];`

---

<sup>1</sup> SQL Server Statistic docs

<sup>2</sup> Store procedure `sp_updatestats`

<sup>3</sup> Update statistics single table

## FEATURES

Firely Server offers many features as defined in the FHIR Specification and beyond.

### 7.1 FHIR RESTful API

Firely Server supports most of the features in the [FHIR RESTful API](#).

#### 7.1.1 FHIR Versions

All the operations below can be called for FHIR STU3 or FHIR R4. Firely Server supports the `fhirVersion` mimetype parameter and `fhir version` endpoint mappings for that purpose. See [Multiple versions of FHIR](#) for more information.

#### 7.1.2 Create, read, update, patch, delete

These five operations to manage the contents of the Firely Server, commonly referenced by the acronym CRUD, are implemented as per the specification. Patch is implemented as [FHIR Patch](#), as this is the most versatile one. This includes version-read and the conditional variations. Only a few limitations apply.

Firely Server enables create-on-update: If you request an update and no resource exists for the given id, the provided resource will be created under the provided id.

Firely Server can reject a resource based on [Validating incoming resources](#).

#### Configuration

A conditional delete interaction may match multiple resources. You can configure the server to delete all matches, or reject the operation (effectively only allowing single matches to be deleted). Allowing multiple deletes requires support for transactions on the database (MongoDb, SQL Server or SQLite). If you allow for multiple deletes, you have to specify a maximum number of resources that can be deleted at once, to save you from accidentally deleting too many resources.

```
"FhirCapabilities": {
  "ConditionalDeleteOptions": {
    "ConditionalDeleteType": "Single", // Single or Multiple,
    "ConditionalDeleteMaxItems": 1
  }
}
```

### Limitations on CRUD

1. Simultaneous conditional creates and updates are not entirely transactional safe:
  - Two conditional updates may both result in a `create`, although the result of one may be a match to the other.
  - Two conditional creates may both succeed, although the result of one may be a match to the other.
  - A conditional create and a simultaneous conditional update may both result in a `create`, although the result of one may be a match to the other.
2. Parameter `_pretty` is not yet supported.
3. XML Patch and JSON Patch are not supported.

### 7.1.3 Versioning

Firely Server keeps a full version history of every resource, including the resources on the *Firely Server Administration API*.

### 7.1.4 Search

Search is supported as per the specification, with a few *Limitations on search*.

In the default configuration the SearchParameters from the *FHIR specification* are available. But Firely Server also allows *Custom Search Parameters*.

Chaining and reverse chaining is fully supported.

Quantity search on UCUM quantities automatically converts units to a canonical form. This means you can have kg in an Observation and search by lbs, or vice versa.

*Compartment Search* is supported.

**Warning:** Queries that request resource types not included in the current compartment's `CompartmentDefinition` will yield default search results. Example: Searching for Practitioner resources within a Patient's compartment will return all Practitioner resources, including the ones not linked to the patient.

Firely Server also supports `_include:iterate` and `_revinclude:iterate`, as well as its STU3 counterparts `_include:recurse` and `_revinclude:recurse`. See *the specification* for the definition of those. You can configure the maximum level of recursion:

```
"FhirCapabilities": {
  "SearchOptions": {
    "MaximumIncludeIterationDepth": 1
  }
},
```

When searching with the `:exact` modifier the server handles *grapheme clusters*.

## Sorting

`_sort` is implemented for searchparameters of types:

- string
- number
- uri
- reference
- datetime
- token

for the all supported repositories.

How is sort evaluated?

- A searchparameter may be indexed with multiple values for a single resource. E.g. Patient.name for Angelina Jolie would have name=Angelina and name=Jolie. And George Clooney: name=George and name=Clooney. As the FHIR Specification phrases it: “In this case, the sort is based on the item in the set of multiple parameters that comes earliest in the specified sort order when ordering the returned resources.” Here is an example of how Firely Server evaluates this.

- In ascending order: Patient?\_sort=name

Name values	Asc. per resource	Asc. resources
Angelina	Angelina	<i>Angelina Jolie</i>
Jolie	Jolie	
George	Clooney	<i>George Clooney</i>
Clooney	George	

- Now in descending order: Patient?\_sort=-name

Name values	Desc. per resource	Desc. resources
Angelina	Jolie	<i>Angelina Jolie</i>
Jolie	Angelina	
George	George	<i>George Clooney</i>
Clooney	Clooney	

- The searchparameter to sort on may not be indexed at all for some of the resources in the resultset. E.g. a Patient without any identifier will not be indexed for Patient.identifier. Resources not having that parameter always end up last (both in ascending and descending order). This is similar to the ‘nulls last’ option in some SQL languages.
- Token parameters are sorted only on their code element. The system element is ignored in the sorting.
- Firely Server uses the default collation as configured on the database server. This collation defines the ordering of characters.
- Sorting on `_score` is not supported.

### Limitations on search

The following parameters and options are not yet supported:

1. `_text`
2. `_content`
3. `_query`
4. `_containedType`
5. `_filter`
6. `Location.near` (geo matching is not supported)
7. `:approx` modifier on a quantity `SearchParameter`
8. `:text` modifier on a string `SearchParameter`
9. `:above`, `:below`, `:in` and `:not-in` modifiers on a token `SearchParameter`, `above` and `below` are also not supported for [Mime Types](#).
10. `:above`, `:below` modifiers on a reference `SearchParameter` (only valid on a [strict hierarchy](#))
11. `_include` and `_revinclude` will match the current version of the referenced resources, also if the reference is versioned.
12. `_pretty`
13. Implicit ranges are supported on dates, datetimes and quantities with a UCUM unit. But not on other quantities and number parameters.
14. Search parameter arguments in exponential form (e.g. `1.8e2`).
15. `_total=estimate`, only `none` and `accurate` are supported.

Furthermore:

1. Paging is supported, but it is not isolated from intermediate changes to resources.

### 7.1.5 History

History is supported as described in the specification, on the system, type and instance level. The `_since` and `_count` parameters are also supported.

### Configuration

```
"HistoryOptions": {  
  "MaxReturnedResults": 100  
}
```

If a `_history` call would result in more than `MaxReturnedResults`, Firely Server asks the user to be more specific. Use this to avoid overloading the server or the connection.

## Limitations on history

1. `_at` parameter is not yet supported.
2. Paging is supported, but it is not isolated from intermediate changes to resources.

### 7.1.6 Batch

Batch is fully supported on the usual endpoint. You can limit the number of entries accepted in a single batch. See [Batch and transaction](#).

Note that batches are not supported in the `/administration` endpoint.

### 7.1.7 Transaction

Transactions are supported, but with the following limitation:

1. The `/administration` endpoint does not support transactions.

You can limit the number of entries accepted in a single transaction. See [Batch and transaction](#).

### 7.1.8 Capabilities

On the Capabilities interaction (`<firely-server-endpoint>/metadata`) Firely Server returns a `CapabilityStatement` that is built dynamically from the supported `ResourceTypes`, `SearchParameters` and interactions. E.g. if you [Configure Search Parameters](#), the `SearchParameters` that are actually loaded appear in the `CapabilityStatement`.

### 7.1.9 Document endpoint

Firely Server supports submitting [FHIR document bundles](#) to the base endpoint of the server. The current version of Firely Server will only extract the unstructured part of the document, i.e. the narrative of the document bundle. The submission of the document will return a `DocumentReference` containing an attachment linking to a `Binary` resource containing the original narrative. Please note that only the top-level narrative will be extracted. No section narrative will be handled. Updates to narratives from documents with the same document identifier will result in an `Update` interaction on the `DocumentReference`.

Please make sure that `Vonk.Plugin.DocumentHandling.DocumentHandlingConfiguration` is enabled in the pipeline options to use this feature.

### 7.1.10 Not supported interactions

These interactions are not yet supported by Firely Server:

1. `HEAD`

Besides that, Firely Server does not yet return the date header as specified in [HTTP return values](#)

## 7.2 Bulk Data Export

Firely Server provides the option to export resources with the Bulk Data Export Service. The Bulk Data Export Service enables the \$export operation from the Fhir specification. Read more about the [\\$export request flow](#)

**Note:** To use Bulk Data Export you have to configure either *SQL Server* or *MongoDB* for the data database. Or you can implement it as part of a *Facade*.

The Administration database can be configured to any of the three supported databases.

### 7.2.1 Appsettings

To start using the Bulk Data Export Service (BDE) you will first have to add the relevant plugins (Vonk.Plugin.BulkDataExport.[Level]BulkDataExportConfiguration) to the PipelineOptions in the appsettings. In the example below we have enabled all three levels: Patient, Group and System.

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",
        "Vonk.Fhir.R4",
        //"Vonk.Fhir.R5",
        "Vonk.Repository.Sql.SqlVonkConfiguration",
        "Vonk.Repository.Sqlite.SqliteVonkConfiguration",
        "Vonk.Repository.MongoDb.MongoDbVonkConfiguration",
        "Vonk.Repository.Memory.MemoryVonkConfiguration",
        "Vonk.Subscriptions",
        "Vonk.Smart",
        "Vonk.UI.Demo",
        "Vonk.Plugin.DocumentOperation.DocumentOperationConfiguration",
        "Vonk.Plugin.ConvertOperation.ConvertOperationConfiguration",
        "Vonk.Plugin.BinaryWrapper",
        "Vonk.Plugin.Audit",
        "Vonk.Plugins.TerminologyIntegration",
        "Vonk.Plugin.BulkDataExport.SystemBulkDataExportConfiguration",
        "Vonk.Plugin.BulkDataExport.GroupBulkDataExportConfiguration",
        "Vonk.Plugin.BulkDataExport.PatientBulkDataExportConfiguration",
      ],
      "Exclude": [
        "Vonk.Subscriptions.Administration"
      ]
    }
  ],
  ...etc...
}
```

Bulk Data Export Service works as an asynchronous operation. To store the all operation-related information, it is necessary to enable a “Task Repository” on the admin database. Please enable the relevant “Vonk.Repository.[database-type].[database-type]TaskConfiguration” in the administration pipeline options, depending on the database type you use for the admin database. All supported databases can be used as a task repository. In the example below we have enabled the task repository for SQLite: “Vonk.Repository.Sqlite.SqliteTaskConfiguration”.



```

"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/administration",
      "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",
        "Vonk.Fhir.R4",
        //"Vonk.Fhir.R5",
        //"Vonk.Repository.Sql.SqlTaskConfiguration",
        //"Vonk.Repository.Sql.SqlAdministrationConfiguration",
        "Vonk.Repository.Sql.Raw.KAdminSearchConfiguration",
        "Vonk.Repository.Sqlite.SqliteTaskConfiguration",
        "Vonk.Repository.Sqlite.SqliteAdministrationConfiguration",
        //"Vonk.Repository.MongoDb.MongoDbTaskConfiguration",
        "Vonk.Repository.MongoDb.MongoDbAdminConfiguration",
        "Vonk.Repository.Memory.MemoryAdministrationConfiguration",
        "Vonk.Subscriptions.Administration",
        "Vonk.Plugins.Terminology",
        "Vonk.Administration",
        "Vonk.Plugin.BinaryWrapper"
      ],
      "Exclude": [
        "Vonk.Core.Operations"
      ],
      ...etc...
    }
  ]
}

```

BDE introduces two new parts to the appsettings, namely TaskFileManagement and BulkDataExport.

```

"TaskFileManagement": {
  "StoragePath": "./taskfiles"
},
"BulkDataExport": {
  "RepeatPeriod" : 60000, //ms
  "AdditionalResources": [ "Organization", "Location", "Substance", "Device",
    ↪ "BodyStructure", "Medication", "Coverage" ]
},

```

In StoragePath you can configure the folder where the exported files will be saved to. Make sure the server has write access to this folder.

In RepeatPeriod you can configure the polling interval (in milliseconds) for checking the Task queue for a new export task.

A patient-based or group-based Bulk Data Export returns resources based on the Patient compartment definition (<https://www.hl7.org/fhir/compartimentdefinition-patient.html>). These resources may reference resources outside the compartment as well, such as a Practitioner who is the performer of a Procedure. Using the *AdditionalResources*-setting, you can determine which types of referenced resources are exported in addition to the compartment resources.

### 7.2.2 \$export

There are three different levels for which the \$export operation can be called:

#### System

**url:** [firely-server-base]/\$export

This will create a system level export task, exporting all resources in the Firely Server database to a .ndjson file per resourcetype.

#### Patient

**url:** [firely-server-base]/Patient/\$export

This will create a type level export task, exporting all resources included in the Patient Compartment in the Firely Server database to an .ndjson file per resourcetype.

#### Group

**url:** [firely-server-base]/Group/<group-id>/\$export

This will create an instance level export task. For each Patient in the Group, the task will export all resources included in the Patient Compartment in the Firely Server database to an .ndjson file per resourcetype.

---

**Note:** For now we only support inclusion in a Group through Group.member.

---

Making an \$export request will create a new task in the database with status “Queued”. The request should return an absolute **\$exportstatus** URL in the Content-Location header and the OperationOutcome in the body.

### 7.2.3 \$exportstatus

The \$export request should return the \$exportstatus url for your export task. This url can be used to request the current status of the task through a GET request, or to cancel the task through a DELETE request.

There are six possible status options:

1. Queued
  2. Active
  3. Complete
  4. Failed
  5. CancellationRequested
  6. Cancelled
- If a task is Queued or Active, GET \$exportstatus will return the status in the X-Progress header
  - If a task is Complete, GET \$exportstatus will return the results with a **\$exportfilerequest** url per exported .ndjson file. This url can be used to retrieve the files per resourcetype. If there were any problems with parts of the export, an url for the generated OperationOutcome resources can be found in the error section of the result.
  - If a task is Failed, GET \$exportstatus will return HTTP Statuscode 500 with an OperationOutcome.

- If a task is on status CancellationRequested or Cancelled, GET \$exportstatus will return HTTP Statuscode 410 (Gone).

## 7.2.4 \$exportfilerequest

If a task has the Complete status, the GET \$exportstatus request should return one or more \$exportfilerequest urls. Performing a GET request on this \$exportfilerequest url returns a body of FHIR resources in newline delimited json (ndjson).

**Note:** The Accept header for this request has to be:

```
application/fhir+ndjson
```

## 7.2.5 Facade

We support BDE for a facade. As always with a facade implementation, the parts dealing with the underlying proprietary datastore need to be implemented by you. Below you find an overview of the relevant steps for implementing BDE for a facade.

Export level	Area	Setting	Action
All	PipelineOptions for the administration endpoint	“Vonk.Repository.[database-type].[database-type]TaskConfiguration”	Enable for relevant administration database type
All	SupportedInteractions.WholeSystemInteractions	\$exportstatus	Enable
All	SupportedInteractions.WholeSystemInteractions	\$exportfilerequest	Enable
All	Facade plugin	IBulkDataExportSnapshotRepository	Implement
Patient	PipelineOptions for the (root) endpoint	“Vonk.Plugin.BulkDataExport.PatientBulkDataExportConfiguration”	Enable
Patient	SupportedInteractions.TypeLevelInteractions	\$export	Enable
Patient	Facade plugin	IPatientBulkDataExportRepository	Implement
Group	PipelineOptions for the (root) endpoint	“Vonk.Plugin.BulkDataExport.GroupBulkDataExportConfiguration”	Enable
Group	SupportedInteractions.InstanceLevelInteractions	\$export	Enable
Group	Facade plugin	IGroupBulkDataExportRepository	Implement
System	PipelineOptions for the (root) endpoint	“Vonk.Plugin.BulkDataExport.SystemBulkDataExportConfiguration”	Enable
System	SupportedInteractions.SystemLevelInteractions	\$export	Enable
System	Facade plugin	ISystemBulkDataExportRepository	Implement

**Note:** The interfaces below can be found in Vonk.Core version 4.7.0 and higher.

### ISystemBulkDataExportRepository

The class implementing this interface is responsible for creating (and eventually deleting) a snapshot of the relevant data. This snapshot will be used at a later time for retrieving the data, mapping it to FHIR and writing the resources to the output files. How you store this snapshot is up to you.

**Attention:** The current implementation of the Bulk Data Export plugin for facades does not trigger `ISystemBulkDataExportRepository.DeleteSnapshot(string taskId)`. This will be resolved in the upcoming release of Firely Server.

### IPatientBulkDataExportRepository

Used when performing a Patient level export. It should retrieve the snapshot, use this to obtain the relevant data from the proprietary datastore and transform this to FHIR resources. Only data directly associated with the relevant Patient resources should be returned.

### IGroupBulkDataExportRepository

Used when performing a Group level export. It should retrieve the snapshot, use this to obtain the relevant data from the proprietary datastore and transform this to FHIR resources.

### ISystemBulkDataExportRepository

Used when performing a System level export. It should retrieve the snapshot, use this to obtain the relevant data from the proprietary datastore and transform this to FHIR resources.

## 7.3 Patient \$everything

### 7.3.1 Description

This plugin implements the Patient \$everything operation, as described in <https://www.hl7.org/fhir/operation-patient-everything.html>. This operation returns all resources associated with a single patient. The resources returned are determined by the Patient compartment, defined in <https://www.hl7.org/fhir/compartmentsdefinition-patient.html>. Currently, the functionality is only available if you use SQL server for your data.

GET <base-url>/Patient/<patient-id>/\$everything  
Accept: <any supported FHIR media type>

Optional parameters:

- `_since`: Get only resources changed since this moment
- `_type`: Limit the returned resource types to only the types in this list

Please note that other defined operation parameters have not been implemented (yet).

So if you would want to fetch only Patient 1 and their Observations, changed since the 1st of January, 2021 in FHIR JSON format, you would use:

```
GET <base-url>/Patient/1/$everything?_type=Patient,Observation&_since=2021-01-01
Accept: application/fhir+json
```

## 7.3.2 Configuration

Many resources in the Patient compartment reference resources outside the compartment. For example: An Observation might have a performer which is a Practitioner. As Practitioner itself is not in the Patient compartment, the resource would normally not be returned. But using a setting you can control which additional resource types are returned if they are referenced by any of the resources you requested.

```
"PatientEverythingOperation": {
  "AdditionalResources": [ "Organization", "Location", "Substance", "Device",
↪ "Medication" ]
}
```

This is the default value for the setting. As you can see, Practitioner is not included by default out of privacy considerations but you can change that by overriding the setting.

**Note:** Device was added as an additional resource as it includes a reference to a patient but is not listed in the patient's compartment yet. As soon as the specification is updated, it will be removed from the appsettings and returned by default.

To include the plugin in your pipeline, add the following extra Include:

```
"PipelineOptions": {
  "Branches": [
    {
      "Path": "/",
      "Include": [
        ...
        "Vonk.Plugin.PatientEverything"
      ]
    },
    ...
  ]
}
```

## 7.3.3 License

The Patient \$everything operation is licensed. To use it, you may need to renew your license.

## 7.4 \$lastN Observations

### 7.4.1 Description

This plugin implements the Observation \$lastn operation, as described in <https://www.hl7.org/fhir/observation-operation-lastn.html>. This operation returns the most recent or last n=number of observations for a patient within a given category (e.g. vital-signs, laboratory, etc.).

```
GET <base-url>/Observation/$lastn
Accept: <any supported FHIR media type>
```

Required parameters:

- **patient / subject**: a reference to the Patient;
- **category**: a category to get the most recent observations from.

Optional parameters:

- **max**: maximum number of Observations to return from each group;
- other search parameters defined for the Observation resource.

### 7.4.2 Appsettings

To start using the \$lastn operation you will first have to add the plugin `Vonk.Plugin.LastN` to the `PipelineOptions` in the appsettings. The plugin depends on `Vonk.Repository.Sql.Raw.KSearchConfiguration` and `Vonk.Repository.MongoDb.MongoDbVonkConfiguration`, which, therefore, should also be included in the pipeline if you use SQL Server or MongoDB respectively.

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        ...
        "Vonk.Repository.Sql.Raw.KSearchConfiguration", // If SQL Server is used
        "Vonk.Repository.MongoDb.MongoDbVonkConfiguration", // If MongoDB is used
        "Vonk.Plugin.LastN",
      ],
      "Exclude": [
        ...
      ]
    }, ...etc...
  ]
}
```

### 7.4.3 Examples

The examples below use a predefined set of resources. You can add those resources to your instance of the server to reproduce the examples. To do it, please execute the following transaction bundle: [download](#). The transaction bundle contains a patient with id = **examplePatient** and a list of Observations for that patient.

**Fetch the last 3 results for all vitals for a patient**

Request:

```
GET <base-url>/Observation/$lastn?max=3&patient=examplePatient&category=vital-signs
Accept: application/fhir+json
```

Response:

```
{
  "resourceType": "Bundle",
  "type": "searchset",
  "timestamp": "2021-09-14T11:52:58.450+04:00",
  "meta": {
    "lastUpdated": "2021-09-14T11:52:58.450+04:00",
    "versionId": "8f3c6c0a-37d7-4fde-b80b-57e72e655fe7"
  },
  "entry": [
    {
      "fullUrl": "<base-url>/Observation/bac874a3-8d89-bd0c-9157-f2fb153d6642",
      "search": {
        "mode": "match"
      },
      "resource": {
        "resourceType": "Observation",
        "id": "bac874a3-8d89-bd0c-9157-f2fb153d6642",
        "status": "final",
        "category": [
          {
            "coding": [
              {
                "system": "http://terminology.hl7.org/CodeSystem/observation-category",
                "code": "vital-signs",
                "display": "vital-signs"
              }
            ]
          }
        ],
        "code": {
          "coding": [
            {
              "system": "http://loinc.org",
              "code": "2708-6",
              "display": "Oxygen saturation in Arterial blood"
            },
            {
              "system": "http://loinc.org",
              "code": "59408-5",
              "display": "Oxygen saturation in Arterial blood by Pulse oximetry"
            }
          ],
          "text": "Oxygen saturation in Arterial blood"
        },
        "subject": {
```

(continues on next page)

(continued from previous page)

```

        "reference": "<base-url>/Patient/examplePatient"
    },
    "effectiveDateTime": "2020-03-03T01:58:48+04:00",
    // ...
  }
},
{
  /* Entry Observation with code [Body Weight(29463-7)] from 2020-03-03 */},
  /* Entry Observation with code [Body Weight(29463-7)] from 2019-04-22 */},
  /* Entry Observation with code [Body Weight(29463-7)] from 2016-04-18 */},
  /* Entry Observation with code [Body Mass Index(39156-5)] from 2019-04-22 */},
  /* Entry Observation with code [Body Mass Index(39156-5)] from 2016-04-18 */},
  /* Entry Observation with code [Body Mass Index(39156-5)] from 2013-04-15 */},
  /* Entry Observation with code [Pain severity - 0-10 verbal numeric rating [Score] -
  ↳ Reported(72514-3)] from 2019-04-22 */},
  /* Entry Observation with code [Pain severity - 0-10 verbal numeric rating [Score] -
  ↳ Reported(72514-3)] from 2016-04-18 */},
  /* Entry Observation with code [Pain severity - 0-10 verbal numeric rating [Score] -
  ↳ Reported(72514-3)] from 2013-04-15 */},
  /* Entry Observation with code [Body Height(8302-2)] from 2019-04-22 */},
  /* Entry Observation with code [Body Height(8302-2)] from 2016-04-18 */},
  /* Entry Observation with code [Body Height(8302-2)] from 2013-04-15 */},
  /* Entry Observation with code [Body temperature(8310-5)], [Oral temperature(8331-
  ↳ 1)] from 2020-03-03 */},
  /* Entry Observation with code [Blood Pressure(85354-9)] from 2020-03-03 */},
  /* Entry Observation with code [Blood Pressure(85354-9)] from 2019-04-22 */},
  /* Entry Observation with code [Blood Pressure(85354-9)] from 2016-04-18 */},
  /* Entry Observation with code [Heart rate(8867-4)] from 2020-03-03 */},
  /* Entry Observation with code [Heart rate(8867-4)] from 2019-04-22 */},
  /* Entry Observation with code [Heart rate(8867-4)] from 2016-04-18 */},
  /* Entry Observation with code [Respiratory rate(9279-1)] from 2020-03-03 */},
  /* Entry Observation with code [Respiratory rate(9279-1)] from 2019-04-22 */},
  /* Entry Observation with code [Respiratory rate(9279-1)] from 2016-04-18 */}
],
"total": 23,
"link": [
  {
    "relation": "self",
    "url": "<base-url>/Observation/$lastn?max=3&patient=examplePatient&category=vital-
    ↳ signs&_count=23&_skip=0"
  }
],
"id": "6d6571c3-e6e0-461e-803f-c044c442191c"
}

```



## Fetch the last laboratory results for a patient

## Request

```
GET <base-url>/Observation/$lastn?patient=examplePatient&category=laboratory
Accept: application/fhir+json
```

## Response

```
{
  "resourceType": "Bundle",
  "type": "searchset",
  "timestamp": "2021-09-14T12:28:40.943+04:00",
  "meta": {
    "lastUpdated": "2021-09-14T12:28:40.943+04:00",
    "versionId": "748c3f1e-1199-44b8-a3c7-f06f1b1b6b49"
  },
  "entry": [
    { /* Entry Observation with code [Microalbumin Creatinine Ratio(14959-1)] from 2019-
    ↪04-22 */},
    { /* Entry Observation with code [Low Density Lipoprotein Cholesterol(18262-6)] from
    ↪2019-04-22 */},
    { /* Entry Observation with code [Carbon Dioxide(20565-8)] from 2019-04-22 */},
    { /* Entry Observation with code [Chloride(2069-3)] from 2019-04-22 */},
    { /* Entry Observation with code [High Density Lipoprotein Cholesterol(2085-9)] from
    ↪2019-04-22 */},
    { /* Entry Observation with code [Total Cholesterol(2093-3)] from 2019-04-22 */},
    { /* Entry Observation with code [Erythrocyte distribution width [Entitic volume] by
    ↪Automated count(21000-5)] from 2016-04-18 */},
    { /* Entry Observation with code [Glucose(2339-0)] from 2019-04-22 */},
    { /* Entry Observation with code [Triglycerides(2571-8)] from 2019-04-22 */},
    { /* Entry Observation with code [Sodium(2947-0)] from 2019-04-22 */},
    { /* Entry Observation with code [Platelet distribution width [Entitic volume] in
    ↪Blood by Automated count(32207-3)] from 2016-04-18 */},
    { /* Entry Observation with code [Platelet mean volume [Entitic volume] in Blood by
    ↪Automated count(32623-1)] from 2016-04-18 */},
    { /* Entry Observation with code [Estimated Glomerular Filtration Rate(33914-3)] from
    ↪2019-04-22 */},
    { /* Entry Observation with code [Creatinine(38483-4)] from 2019-04-22 */},
    { /* Entry Observation with code [Hematocrit [Volume Fraction] of Blood by Automated
    ↪count(4544-3)] from 2016-04-18 */},
    { /* Entry Observation with code [Hemoglobin A1c/Hemoglobin.total in Blood(4548-4)]
    ↪from 2019-04-22 */},
    { /* Entry Observation with code [Calcium(49765-1)] from 2019-04-22 */},
    { /* Entry Observation with code [Potassium(6298-4)] from 2019-04-22 */},
    { /* Entry Observation with code [Urea Nitrogen(6299-2)] from 2019-04-22 */},
    { /* Entry Observation with code [Leukocytes [# /volume] in Blood by Automated
    ↪count(6690-2)] from 2016-04-18 */},
    { /* Entry Observation with code [Hemoglobin [Mass/volume] in Blood(718-7)] from 2016-
    ↪04-18 */},
    { /* Entry Observation with code [Platelets [# /volume] in Blood by Automated
    ↪count(777-3)] from 2016-04-18 */},
    { /* Entry Observation with code [MCH [Entitic mass] by Automated count(785-6)] from
    ↪2016-04-18 */},
```

(continues on next page)

(continued from previous page)

```

    { /* Entry Observation with code [MCHC [Mass/volume] by Automated count(786-4)] from_
↳ 2016-04-18 */},
    { /* Entry Observation with code [MCV [Entitic volume] by Automated count(787-2)]_
↳ from 2016-04-18 */},
    { /* Entry Observation with code [Erythrocytes [#/volume] in Blood by Automated_
↳ count(789-8)] from 2016-04-18 */},
    { /* Entry Observation with code [Rhinovirus RNA [Presence] in Respiratory specimen_
↳ by NAA with probe detection(92130-4)] from 2020-03-03 */},
    { /* Entry Observation with code [Respiratory syncytial virus RNA [Presence] in_
↳ Respiratory specimen by NAA with probe detection(92131-2)] from 2020-03-03 */},
    { /* Entry Observation with code [Human metapneumovirus RNA [Presence] in Respiratory_
↳ specimen by NAA with probe detection(92134-6)] from 2020-03-03 */},
    { /* Entry Observation with code [Parainfluenza virus 3 RNA [Presence] in Respiratory_
↳ specimen by NAA with probe detection(92138-7)] from 2020-03-03 */},
    { /* Entry Observation with code [Parainfluenza virus 2 RNA [Presence] in Respiratory_
↳ specimen by NAA with probe detection(92139-5)] from 2020-03-03 */},
    { /* Entry Observation with code [Parainfluenza virus 1 RNA [Presence] in Respiratory_
↳ specimen by NAA with probe detection(92140-3)] from 2020-03-03 */},
    { /* Entry Observation with code [Influenza virus B RNA [Presence] in Respiratory_
↳ specimen by NAA with probe detection(92141-1)] from 2020-03-03 */},
    { /* Entry Observation with code [Influenza virus A RNA [Presence] in Respiratory_
↳ specimen by NAA with probe detection(92142-9)] from 2020-03-03 */},
    { /* Entry Observation with code [Adenovirus A+B+C+D+E DNA [Presence] in Respiratory_
↳ specimen by NAA with probe detection(94040-3)] from 2020-03-03 */},
    { /* Entry Observation with code [SARS-CoV-2 RNA Pnl Resp NAA+probe(94531-1)] from_
↳ 2020-03-03 */}
  ],
  "total": 36,
  "link": [
    {
      "relation": "self",
      "url": "<base-url>/Observation/$lastn?patient=examplePatient&category=laboratory&_
↳ count=36&_skip=0"
    }
  ],
  "id": "b6521ba6-6235-4221-95cd-e0f25edd77dc"
}

```

### Get the most recent Observations in category vital-signs conducted before January 1, 2015

#### Request

```

GET <base-url>/Observation/$lastn?patient=examplePatient&category=vital-signs&
↳ date=lt2015-01-01
Accept: application/fhir+json

```

#### Response

```

{
  "resourceType": "Bundle",
  "type": "searchset",

```

(continues on next page)

(continued from previous page)

```

"timestamp": "2021-09-14T12:35:32.952+04:00",
"meta": {
  "lastUpdated": "2021-09-14T12:35:32.952+04:00",
  "versionId": "1b88af29-6f90-4a73-8d21-bf4594f45fec"
},
"entry": [
  /* Entry Observation with code [Body Weight(29463-7)] from 2013-04-15 */},
  /* Entry Observation with code [Body Mass Index(39156-5)] from 2013-04-15 */},
  /* Entry Observation with code [Pain severity - 0-10 verbal numeric rating [Score] -
  ↳ Reported(72514-3)] from 2013-04-15 */},
  /* Entry Observation with code [Body Height(8302-2)] from 2013-04-15 */},
  /* Entry Observation with code [Blood Pressure(85354-9)] from 2013-04-15 */},
  /* Entry Observation with code [Heart rate(8867-4)] from 2013-04-15 */},
  /* Entry Observation with code [Respiratory rate(9279-1)] from 2013-04-15 */}
],
"total": 7,
"link": [
  {
    "relation": "self",
    "url": "<base-url>/Observation/$lastn?patient=examplePatient&category=vital-signs&
    ↳ date=lt2015-01-01&_count=7&_skip=0"
  }
],
"id": "b4178262-9bd3-4d9e-b4de-1578cb5d92de"
}

```

### Fetch the last 3 body weight and body height measurements for a patient

#### Request

```

GET <base-url>/Observation/$lastn?max=3&patient=examplePatient&category=vital-signs&
  ↳ code=29463-7,8302-2
Accept: application/fhir+json

```

#### Response

```

{
  "resourceType": "Bundle",
  "type": "searchset",
  "timestamp": "2021-09-14T12:55:06.929+04:00",
  "meta": {
    "lastUpdated": "2021-09-14T12:55:06.929+04:00",
    "versionId": "3dd3bcde-cbfb-4003-98d7-d7c2f3194c8a"
  },
  "entry": [
    /* Entry Observation with code [Body Weight(29463-7)] from 2020-03-03 */},
    /* Entry Observation with code [Body Weight(29463-7)] from 2019-04-22 */},
    /* Entry Observation with code [Body Weight(29463-7)] from 2016-04-18 */},
    /* Entry Observation with code [Body Height(8302-2)] from 2019-04-22 */},
    /* Entry Observation with code [Body Height(8302-2)] from 2016-04-18 */},
    /* Entry Observation with code [Body Height(8302-2)] from 2013-04-15 */}
  ]
}

```

(continues on next page)

(continued from previous page)

```
],
"total": 6,
"link": [
  {
    "relation": "self",
    "url": "<base-url>/Observation/$lastn?max=3&patient=examplePatient&category=vital-
↪ signs&code=29463-7,8302-2&_count=6&_skip=0"
  }
],
"id": "49ee0b4b-00bd-40b7-8cb5-96a0e0892380"
}
```

## 7.4.4 License

The \$lastn operation is part of the core Firely Server functionality. However, to use it, you may need to request an updated license from Firely. You can use your current license file if it contains <http://fire.ly/vonk/plugins/lastn>.

## 7.5 Custom Operations

### 7.5.1 Validation

Firely Server can validate a resource against a profile as defined in the [\\$validate](#) operation.

You can call validate on three levels:

1. *Validate on the system level*
2. *Validate on the ResourceType level*
3. *Validate an instance from the database*

Besides that you can configure Firely Server to validate every incoming resource and even filter on specific profiles. See the section on [Validating incoming resources](#). In all cases, the *Precondition* is that Firely Server must have access to all relevant StructureDefinitions.

Validation has some *Limitations*.

---

**Note:** The very first validation call will take a considerable amount of time, typically around 5 seconds. This is because Firely Server maintains a cache of validation information, and on the first call that cache is still empty. Subsequent calls are much faster.

---

## Validate on the system level

```
POST <firely_server_endpoint>/$validate[?profile=<canonical-url-of-structuredefinition>]
```

There are two ways of calling \$validate:

1. With a Resource or a Bundle of resources as body, and optionally, a profile parameter on the url.
2. With a Parameters resource as body, having
  - a parameter element with the Resource to validate in the resource parameter;
  - (optionally) the profile to validate against in the profile parameter

In both cases the request must have a Content-Type header matching the format of the body (`application/fhir+json` or `application/fhir+xml`).

If you do not specify a profile parameter, Firely Server will validate the Resource against any profiles mentioned in `meta.profile` as well as the base profile from the FHIR Specification.

If you call \$validate on the system level, Firely Server will make no assumptions about the ResourceType of the Resource to validate.

## Validate on the ResourceType level

```
POST <firely_server_endpoint>/<resourcetype>/$validate[?profile=<canonical-url-of-structuredefinition>]
```

You can call \$validate in the same two ways as with *Validate on the system level*.

If you call \$validate on the ResourceType level, Firely Server will check whether the Resource to validate is of the same `<resourcetype>` as provided in the url.

## Validate an instance from the database

```
GET <firely_server_endpoint>/<resourcetype>/<id>/$validate[?profile=<canonical-url-of-structuredefinition>]
```

This time you can only use the (optional) profile parameter on the url to specify a StructureDefinition to validate against.

## Precondition

Firely Server must be aware of all the StructureDefinitions referenced directly via parameter or indirectly by a profile in `meta.profile`. Refer to the *Managing Conformance Resources* for more information.

### Limitations

1. The mode parameter is not yet supported.
2. Implicit ValueSets (ones that use the .filter property) are not supported - create explicit ones instead (without the .filter property).

### 7.5.2 Snapshot generation

Firely Server is capable of generating a snapshot for a StructureDefinition. This operation is not defined in the FHIR Specification.

You can invoke this operation with

```
POST <firely-server-endpoint>/StructureDefinition/$snapshot
```

- The body must contain the StructureDefinition that you want filled with a fresh snapshot. The StructureDefinition may contain an existing snapshot, it will be ignored.
- The Content-Type header must match the format of the body (application/fhir+json or application/fhir+xml)

Firely Server will return the same StructureDefinition, but with the snapshot element (re-)generated.

---

**Note:** The very first call to \$snapshot will take a considerable amount of time, typically around 5 seconds. This is because Firely Server maintains a cache of StructureDefinition information, and on the first call that cache is still empty. Subsequent calls are much faster.

---

### Precondition

Firely Server must be aware of all the other StructureDefinitions that are referred to by the StructureDefinition in the body of the request. Refer to the *Managing Conformance Resources* for more information.

### 7.5.3 \$meta

Firely Server provides an implementation of the \$meta operation as defined in the [FHIR Specification](#).

By default the operation is only enabled on the level of a resource instance. It can also be enabled on the level of a resourcetype or system wide, but the cost of execution will then be high. On sufficient customer demand an optimized implementation is possible.

The FHIR Specification [operations framework](#) allows for the definition of custom operations and defines how to offer them in the [FHIR RESTful API](#). Firely Server offers two custom operations out of the box.

1. *Validation*
2. *Snapshot generation*

## 7.6 Custom Resources

Custom Resources are not formally defined in the FHIR Specification. To Firely Server a Custom Resource is a resource with a definition that is a specialization of `DomainResource`, but that is not in the Core FHIR Specification. Firely Server can handle these, provided it knows about the `StructureDefinition` that defines it. This page explains how to register such a `StructureDefinition` and store custom resources.

**Warning:** Custom Resources are not interoperable. Do not use them for exchanging resources outside your own programming control boundaries.

### 7.6.1 What to use them for?

Firely Server can be used as a platform to build apps on. In these apps, structures arise outside of the FHIR Specification or even the Health domain. Still, it would be useful to also use Firely Server to store, search and version these structures. Note that this is only for internal use in the app.

### 7.6.2 Register the definition

Just like any resourcetype, the definition for a custom resource is formalized as a `StructureDefinition`. Firely Server will recognize it as the definition of a custom resourcetype if and only if:

- base = `DomainResource`
- derivation = specialization
- kind = resource
- abstract = false

This also means that a Logical Model as-is cannot be used as the definition of a custom resourcetype.

Examples of these can be found in the specification: each resourcetype is defined this way. The easiest way to get started is with the definition of `Basic` (in [xml](#) or [json](#)), and adjust that:

1. Choose a name for the type, let's say 'Foo'.
2. Choose a url for the type. In STU3 this has to start with <http://hl7.org/fhir/StructureDefinition/> (constraint sdf-7), so <http://hl7.org/fhir/StructureDefinition/Foo> makes sense. Note that in R4 you are encouraged to use a url in a domain that you control and *not* within hl7.org.
3. Make sure the id, name and type elements align with the name 'Foo'.
4. Adjust the description
5. Make sure all the elements in the differential start with 'Foo.'
6. (Recommended) Store your definition in Simplifier.net for version management, comments and collaboration.

If you have created the `StructureDefinition`, register it in Firely Server using any of the methods mentioned in [Managing Conformance Resources](#). As an example we will issue an update interaction on the Administration API:

```
PUT <base-url>/administration/StructureDefinition/Foo
Content-Type=application/fhir+json; fhirVersion=3.0
```

By using an update we can choose the id and hence the location of this `StructureDefinition`. Firely Server does this by default for all the resourcetypes defined by the specification as well.

### 7.6.3 Use a resource

To test whether you can actually use the endpoint associated with your custom resourcetype, try a search on it: GET `<base-url>/Foo`. This should return an empty bundle. If you get an `OperationOutcome` with the text “Request for not-supported ResourceType(s) Foo”, the registration of the definition is not correct.

---

**Note:** The `CapabilityStatement` will not list the custom definition. This is because the element `CapabilityStatement.rest.resource.type` has a Required binding to the `ResourceType` valueset. And obviously this valueset does not contain our ‘Foo’ resourcetype.

---

Now use your favorite editor to create a resource that conforms to the Foo `StructureDefinition`. And then create it on Firely Server: POST `<base-url>/Foo`.

All the operations on specification-defined resourcetypes are also available for custom resources. You can also use them in a batch or transaction bundle. Custom Resources can also be validated. This also means that *Validating incoming resources* can be used in conjunction with Custom Resources.

### 7.6.4 Search parameters on a custom resource

In Firely Server you can define your own custom search parameters on any type of resource (see *Custom Search Parameters*). This includes Custom Resources. Just use the type name of the Custom Resource in the `SearchParameter.base`.

## 7.7 Terminology services

Firely Server provides support for Terminology operations and validation against terminologies. This is done through a local implementation based on the Administration database (‘Local Terminology Service’) and - if configured - extended with external FHIR Terminology Servers (‘Remote Terminology Services’). The configuration allows you to set preferred services for each `CodeSystem` and `ValueSet`. Firely Server will then transparently select and query either the Local or one of the Remote Terminology Services.

Of the operations listed below the following can be supported by the Local Terminology Service: `$validate-code`, `$expand`, `$lookup`, `$find-matches`. Note that it only supports simple `ValueSets` and `CodeSystems` like the ones part of the FHIR specification. It cannot support complex terminologies like LOINC or SNOMED-CT.

The terminology operations can be invoked for different FHIR versions as specified in *Multiple versions of FHIR*.

### 7.7.1 Terminology Integration

In earlier versions of Firely Server, local terminology services were separately configured from so-called Terminology Integration. These are now merged.



## Operations

### ValueSet \$validate-code

**definition**

<http://www.hl7.org/implement/standards/fhir/valueset-operation-validate-code.html>

**notes**

- Available on the type level `<firely-server-endpoint>/administration/ValueSet/$validate-code` and the instance level `<firely-server-endpoint>/administration/ValueSet/<id>/$validate-code`.
- Only the parameters url, valueSet, valueSetVersion, code, system, display, coding, codeableConcept, abstract are supported.
- The url and valueSetVersion input parameters are only taken into consideration if no valueSet resource was provided in the body. So the valueSet in the body takes priority.
- Both GET and POST interactions are available.

### ValueSet \$expand

**definition**

<http://www.hl7.org/implement/standards/fhir/valueset-operation-expand.html>

**notes**

- Available on the type level `<firely-server-endpoint>/administration/ValueSet/$expand` and the instance level `<firely-server-endpoint>/administration/ValueSet/<id>/$expand`.
- Only the parameters url, valueSet, valueSetVersion and includeDesignations are supported.
- The url and valueSetVersion input parameters are only taken into consideration if no valueSet resource was provided in the body. So the valueSet in the body takes priority.
- Both GET and POST interactions are available.

### CodeSystem \$lookup

**definition**

<http://www.hl7.org/implement/standards/fhir/codesystem-operation-lookup.html>

**notes**

- Available on the type level `<firely-server-endpoint>/administration/CodeSystem/$lookup`.
- Only the parameters code, system, version, coding and date are supported.
- Code & system combination takes priority over the coding parameter.
- Both GET and POST interactions are available.

### CodeSystem \$find-matches / \$compose

#### definition

<http://www.hl7.org/implement/standards/fhir/codesystem-operation-find-matches.html>

#### notes

- Available on the type level <firely-server-endpoint>/administration/CodeSystem/\$find-matches and the instance level <firely-server-endpoint>/administration/CodeSystem/<id>/\$find-matches.
- Only the parameters system, exact, version, property.code and property.value are supported.
- The url and valueSetVersion input parameters are only taken into consideration if no valueSet resource was provided in the body. So the valueSet in the body takes priority.
- Both GET and POST interactions are available.
- \$find-matches was named \$compose in FHIR STU3. The operation is supported with both names.

### CodeSystem \$subsumes

#### definition

<http://www.hl7.org/implement/standards/fhir/codesystem-operation-subsumes.html>

#### notes

- Available on the type level <firely-server-endpoint>/administration/CodeSystem/\$subsumes.
- Only the parameters codeA, codeB, system, and version are supported.
- The system input parameters is only taken into when called on the type level.
- Both GET and POST interactions are available.

### ConceptMap \$closure

#### definition

<http://www.hl7.org/implement/standards/fhir/conceptmap-operation-closure.html>

#### notes

- Available on the system level <firely-server-endpoint>/administration/\$closure.
- This operation is passed on to a Remote Terminology Service supporting it. It supports any parameters that the Remote service supports.
- Only POST interactions are available.

## ConceptMap \$translate

### definition

<http://www.hl7.org/implement/standards/fhir/conceptmap-operation-translate.html>

### notes

- Available on the instance level `<firely-server-endpoint>/administration/ConceptMap/[id]/$translate` and the type level `<firely-server-endpoint>/administration/ConceptMap/$translate`.
- Only the parameters `url`, `conceptMap` (on POST), `conceptMapVersion`, `code`, `system`, `version`, `source`, `target`, `targetsystem` and `reverse` are supported.
- Both GET and POST interactions are available.

## Configuration

### 7.7.2 Pipeline

Make sure to add the `Vonk.Plugins.Terminology` plugin to the `PipelineOptions` in `appsettings` in order to make use of the `TerminologyIntegration` plugin. Additionally, to the “/administration” pipeline, `Vonk.Plugins.Terminology` can be used on the regular FHIR pipeline “/”. Please note that in this case, `CodeSystems` and `ValueSets` are resolved from the Administration repository when executing a terminology operation and the corresponding resource is not provided as part of the request as a parameter.

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [...]
    },
    {
      "Path": "/administration",
      "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",
        "Vonk.Fhir.R4",
        "Vonk.Administration",
        ...
        "Vonk.Plugins.Terminology"
      ],
      "Exclude": [
        "Vonk.Subscriptions.Administration"
      ]
    },
    ...etc...
  ]
}
```

To include or exclude individual operations in the pipeline, see the available plugins under *Terminology*.

Also make sure that the terminology operations are allowed at all in the `SupportedInteractions` section:

```
"SupportedInteractions": {
  "InstanceLevelInteractions": "$validate-code, $expand, $compose, $translate, $subsumes
  ↪",
}
```

(continues on next page)

(continued from previous page)

```

    "TypeLevelInteractions": "$validate-code, $expand, $lookup, $compose, $translate,
    ↪$subsumes",
    "WholeSystemInteractions": "$closure"
  },

```

Lastly, operation on the administration endpoint can be limited to specific IP addresses:

```

"Administration": {
  "Security": {
    "AllowedNetworks": [ "127.0.0.1", "::1" ], // i.e.: [ "127.0.0.1", "::1" (ipv6_
    ↪localhost), "10.1.50.0/24", "10.5.3.0/24", "31.161.91.98" ]
    "OperationsToBeSecured": [ "$validate-code", "$expand", "$compose", "$translate", "
    ↪$subsumes", "$lookup", "$closure" ]
  }
},

```

## 7.7.3 Options

You can enable the integration with one or more external terminology services by setting the required options in the appsettings file. There is a block for the Local Terminology Service and one for each Remote Terminology Service.

For each terminology service you can set the following options:

### Order

The order of the terminology service, or the priority. If multiple Terminology services could be used for a request, Firely Server will use the priority to select a service. Terminology services are arranged in a ascending order: so 1 will be selected over 2.

### PreferredSystem

If a request is directed at a specific code system, Firely Server will choose this terminology server over other available services. A system matches one of the preferred systems if the system starts with the preferred system. So <http://loinc.org> will match any CodeSystem or ValueSet with a canonical that starts with that url.

### SupportedInteractions

The operations supported by the terminology service. Firely Server will only select this service if the operation is in this list. Valid values:

```

"ValueSetValidateCode"
"CodeSystemValidateCode"
"Expand"
"FindMatches" / "Compose"
"Lookup"
"Translate"
"Subsumes"
"Closure"

```

### SupportedInformationModels

The FHIR versions supported by the terminology service. Valid values:

```

"Fhir3.0"
"Fhir4.0"
"Fhir5.0"

```

**Endpoint**

The endpoint url where Firely Server can redirect the requests to.

**Username**

If the terminology service uses Basic Authentication, you can set the required username here.

**Password**

If the terminology service uses Basic Authentication, you can set the required password here.

**MediaType**

Default Media-Type that should be used for serialization of the Parameters resources forwarded to the external terminology service

Notes:

- The Endpoint, Username and Password settings are not valid for the Local Terminology Server, just for the Remote services.
- If a Remote Terminology Service has different endpoints for different FHIR versions, configure each endpoint separately.
- The SupportedInformationModels cannot be broader than the corresponding Fhir.Rx plugins configured in the PipelineOptions.

A sample Terminology section in the appsettings can look like this:

```
"Terminology": {
  "MaxExpansionSize": 650,
  "LocalTerminologyService": {
    "Order": 10,
    "PreferredSystems": [ "http://hl7.org/fhir" ],
    "SupportedInteractions": [ "ValueSetValidateCode", "Expand" ],
    "SupportedInformationModels": [ "Fhir3.0", "Fhir4.0", "Fhir5.0" ]
  },
  "RemoteTerminologyServices": [
    {
      "Order": 20,
      "PreferredSystems": [ "http://snomed.info/sct" ],
      "SupportedInteractions": [ "ValueSetValidateCode", "Expand", "Translate", "Subsumes
↪", "Closure" ],
      "SupportedInformationModels": [ "Fhir4.0" ],
      "Endpoint": "https://r4.ontoserver.csiro.au/fhir/",
      "MediaType": "application/fhir+xml"
    },
    {
      "Order": 30,
      "PreferredSystems": [ "http://loinc.org" ],
      "SupportedInteractions": [ "ValueSetValidateCode", "Expand", "Translate" ],
      "SupportedInformationModels": [ "Fhir3.0", "Fhir4.0" ],
      "Endpoint": "https://fhir.loinc.org/",
      "Username": "",
      "Password": ""
    }
  ]
},
```

This means if you execute a terminology operation request, Firely Server will check whether the request is correct, redirect it to the preferred terminology service and finally return the result.

Additionally to the remote and local terminology services, you can configure the maximum number of concepts that are allowed to be included in a local ValueSet expansion (MaxExpansionSize). ValueSets stored in the local administration database larger than the configured setting will not be expanded, hence they cannot be used for \$validate-code, \$validate or \$expand.

### License

The Terminology plugin itself is licensed with the license token <http://fire.ly/vonk/plugins/terminology>.

When you configure Remote Terminology Services it is your responsibility to check whether you are licensed to use those services.

## 7.8 Custom Search Parameters

### 7.8.1 Configure Search Parameters

You can control which search parameters are known to Firely Server. This is managed in the same way as all the conformance resources, see [Managing Conformance Resources](#).

### 7.8.2 Re-indexing for new or changed SearchParameters

Firely Server extracts values from resources based on the available search parameters upon create or update. This means that if you already had resources in your database before adding a custom search parameter, those resources will not be indexed for that parameter. If you on the other hand removed a previously used search parameter, the index will contain superfluous data.

To fix that, you should re-index (repeat the extraction) for these parameters.

In short, both reindex operations below will:

- Return an Operation Outcome stating that the reindex procedure was started successfully.
- Run the actual reindex asynchronously, using a configured number of threads, thereby using most of the hardware resources available to Firely Server.
- Block any other requests for the duration of the reindex.
- Log progress in the log.

**Caution:** This is a possibly lengthy operation, so use it with care.

- Always try the reindex on a representative (sub)set of your data in a test environment to assess how long the operation may take in the production environment.
- Always make a backup of your data before performing a reindex.

**Warning:** During the re-index operation, all other operations are blocked and responded to with response code '423 - Locked'.

## Reindexing and FHIR versions

Reindexing is also controlled by the `fhirVersion` parameter (see [Multiple versions of FHIR](#)) in the Accept header or the version-mapped endpoint. It will then reindex only for SearchParameters and resources *in that FHIR version*. So for a full reindex of everything you may need to issue the command twice, once for each `fhirVersion`.

### Rebuild the whole search index

This is only needed if we changed something very significant to the way Firely Server searches, like

- The way values are extracted for all or many searchparameters.
- The structure in which Firely Server stores the search index.

To re-index all resources for all search parameters, use:

```
POST http(s)://<firely-server-endpoint>/administration/reindex/all
Accept=application/fhir+json (or xml); fhirVersion=3.0 (or 4.0)
```

This will delete any previously indexed data and extract it again from the resources.

### Rebuild the search index for specific searchparameters

This is needed if:

- The definition (usually the expression) of a searchparameter has changed.
- A searchparameter was added.
- A searchparameter was removed and you want the search index to be tidy and not have this parameter in it anymore.

To re-index all resources for certain search parameters, use:

```
POST http(s)://<firely-server-endpoint>administration/reindex/searchparameters
Accept=application/fhir+json (or xml); fhirVersion=3.0 (or 4.0)
```

In the body of the POST, you put the name of the search parameters to actually re-index as form parameters:

```
include=Patient.name,Observation.code
exclude=Organization.name
```

`include` means that resources will be re-indexed only for those search parameters. You use this if you added or changed one or few search parameters.

`exclude` means that any existing index data for those search parameters will be erased. You use this when you removed a search parameter.

Remember to adjust the Content-Type header: `application/x-www-form-urlencoded`.

If you are *not permitted* to perform the reindex, Firely Server will return statuscode 403.

## Re-index Configuration

Firely Server will not re-index the resources in the database all at once, but in batches. The re-index operation will process all batches until all resources are re-indexed. You can control the size of the batches in the [Firely Server settings](#). Besides that you can also control how many threads run in parallel to speed up the reindex process. The configured value is a maximum, since Firely Server will also be limited by the available computing resources.

```
"ReindexOptions": {  
  "BatchSize": 100,  
  "MaxDegreeOfParallelism": 10  
},
```

Use any integer value  $\geq 1$ .

## 7.8.3 Limitations

Every search parameter has to have either:

- a valid FhirPath in it's Expression property, or
- be a Composite search parameter and specify at least one component.

## 7.9 Errata to the specification

The FHIR Specification is good, but not perfect. Some of the SearchParameters have errors. If we find these errors, we report them in the issue tracking system of HL7. But it takes time until the fix is applied to the specification. In the meantime Firely Server provides you with updated versions of the resources that have errors, so you can use them already while we await the fixes in the specification.

These corrections come with the Firely Server installation, in the files:

- `errataFhir3.0.zip`, with corrections to the STU3 version of the Specification
- `errataFhir4.0.zip`, with corrections to the R4 version of the Specification

These files are imported automatically during startup, as are other conformance resources, see [Managing Conformance Resources](#).

Currently the `errata.zip` file contains the following corrections:

### clinical-patient

This parameter incorrectly specified that both Patient and Group were target resource types for the patient search parameter. For DeviceUseStatement-patient this was correct, so we created a separate file for this parameter, still listing the Group as a valid target type.

### search parameters with FhirPath expression `.as(DateTime)`

Several search parameters had an incorrect FhirPath expression using `.as(DateTime)` instead of `.as(dateTime)`. As a result, Firely Server could not index the fields correctly and searches on the dates would not work. The search parameters that were corrected are: `clinical-date`, `DeviceRequest-event-date`, `Observation-code-value-date`, `Observation-value-date` and `patient-death-date`.

### Resource.<xyz> expressions

The FhirPath library did not support polymorphism yet, so all the search parameters defined with an expression of `Resource.<xyz>` – for example `Resource.meta.lastUpdated` – did not work correctly. We have changed the expression to have just the `<xyz>` part – for example `meta.lastUpdated`.



**StructureDefinition.ext-context (R4 only)**

The FhirPath expression ended on a BackboneElement that cannot be indexed. Changed to the expression `StructureDefinition.context.where(type='element').expression`.

## 7.10 Subscriptions

Subscriptions can be managed in the *Firely Server Administration API*, on the `/administration/Subscription` endpoint. If you post a Subscription to the regular FHIR endpoint, it will be stored but not evaluated. Subscriptions posted to the `/administration` endpoint will be processed and evaluated for each POST/PUT to the server.

Firely Server currently only supports STU3/R4-style Subscriptions.

If you are *not permitted* to access the `/Subscription` endpoint, Firely Server will return statuscode 403.

See [Subscriptions in the specification](#) for more background on Subscriptions.

### 7.10.1 FHIR versions

You POST a Subscription with a `fhirVersion` parameter (see *Multiple versions of FHIR*) or to a version specific endpoint. It will then respond to changes on resources *in that FHIR version*. So if you need a Subscription on both STU3 and R4 resources, POST that Subscription for both FHIR versions.

### 7.10.2 Channels

According to the FHIR specification, a channel defines how a FHIR server notifies other systems when resources get created or updated. The specification describes several channel types. Currently, Firely Server supports only *rest-hook* channel type.

Below is an example of a Subscription resource that uses this channel type.

```
{
  "resourceType": "Subscription",
  "id": "example",
  "status": "requested",
  "contact": [
    {
      "system": "phone",
      "value": "ext 4123"
    }
  ],
  "end": "2025-01-01T00:00:00Z",
  "reason": "Monitor new neonatal function",
  "criteria": "Observation?code=http://loinc.org|1975-2",
  "channel": {
    "type": "rest-hook",
    "endpoint": "https://my-subscription-endpoint.com",
    "payload": "application/fhir+json",
    "header": [
      "Authorization: Bearer secret-token-abc-123"
    ]
  }
}
```

Once this Subscription resource is posted to Firely Server, the server will be sending notifications to the specified *endpoint* whenever a resource matching the search *criteria* gets created or updated. It is possible to provide *headers* that will be copied over to the notification requests. It may come in handy if the notifications endpoint is secured and the Authorization header must be used. The *payload* option defines the format of the notification payload. The following values can be used:

- *application/fhir+json* and *application/json*
- *application/fhir+xml* and *application/xml*

The payload will contain the created/updated resource in the FHIR format. If the payload option is omitted, the notifications will be sent without the body.

### 7.10.3 Configuration

Firely Server evaluates the active Subscriptions periodically, and in batches (x at a time, until all the active Subscriptions have been evaluated). You can control the period and the batchsize. If an evaluation of a Subscription fails, Firely Server will retry the evaluation periodically for a maximum amount of tries. You can control the retry period and the maximum number of retries.

```
"SubscriptionEvaluatorOptions": {  
  "Enabled" : true  
  "RepeatPeriod": 20000,  
  "SubscriptionBatchSize" : 1,  
  "RetryPeriod": 60000,  
  "MaximumRetries": 3,  
  "SendRestHookAsCreate": false  
},
```

- **Enabled** allows you to quickly enable or disable the evaluation of Subscriptions. Default value is 'false', which implies that Subscription evaluation is also off if this section is left out of the settings.
- **RepeatPeriod** is expressed in milliseconds. In the example above the period is set to 20 seconds, meaning that after a change a subscriber will be notified in at most 20 seconds.
- **SubscriptionBatchSize** is expressed in number of Subscriptions that is retrieved and evaluated at once. Default is 1, but you can set it higher if you have a lot of Subscriptions.
- **RetryPeriod** is expressed in milliseconds. In the example above the period is set to 60 seconds, meaning that Firely Server will retry to send the resources after a minimum of 60 seconds. Retry is included in the normal evaluation process, so the **RetryPeriod** cannot be smaller than **RepeatPeriod**.
- **MaximumRetries** is the maximum amount of times Firely Server will retry to send the resources.
- **SendRestHookAsCreate**: in versions < 3.9.3, Vonk sent RestHook notifications as a create operation using a POST. That was not compliant with the specification that requires an update operation using a PUT. The default value of **false** provides compliant behaviour. Only set it to **true** if you need Firely Servder to keep sending create operations as it did previously.

## 7.11 X-Provenance header

### 7.11.1 Description

The X-Provenance header can be used to add a Provenance resource upon creating or updating another resource. See [Provenance#header](#) for more information.

---

**Note:** In the case of a conditional create, where the resource was already present, the Provenance resource is not created.

---

### Facade implementations

---

**Note:** The provenance header functionality is dependent on transaction services. If transaction support is disabled in your facade implementation, the configuration for the provenance header should be disabled as well.

---

### Linking the Provenance resource

Provenance must point to its target resource through the `target` property (see [Provenance#resource](#)). The `target` will be automatically deduced from the created/updated resource and does not need to be included in the header. The reference will point to the current version (after create/update) of the resource.

### Validation of the Provenance resource

Validation of the Provenance resource follows the settings for validation (see [Validating incoming resources](#)).

---

**Note:** If validation is disabled, an invalid Provenance resource in the header will be ignored and the request will complete normally (if valid).

---

### Limitations

The maximum header length of the web server applies (e.g. IIS 8KB/16KB).

### 7.11.2 Example

```
POST {base-url}/Patient

X-Provenance: { "resourceType": "Provenance", "text": { "status": "generated", "div": "
↪<div>Record of change</div>" }, "recorded": "2022-08-24T11:05:24+02:00", "agent": [ {
↪"who": { "reference": "[mandatory reference]" } } ] }

{
  "resourceType": "Patient",
  "active": true,
```

(continues on next page)

(continued from previous page)

```
"name": [
  {
    "use": "official",
    "family": [
      "Doe"
    ],
    "given": [
      "John"
    ]
  }
],
"gender": "male",
"birthDate": "1974-12-25"
}
```

Update [mandatory reference] to point to a valid resource in your system.

### 7.11.3 License

The X-Provenance header is available for licenses with the Transaction feature enabled.

## 7.12 Reset the database

If you have set up Firely Server as a reference server in a testing environment, it can be useful to reset the database. You would usually do this in combination with *Preloading resources*.

To reset the database, execute:

```
POST http(s)://<firely-server-endpoint>/administration/reset
```

Firely Server will return statuscode 200 if the operation succeeded.

If you are *not permitted* to perform the reset, Firely Server will return statuscode 403.

---

**Note:** On a large database this operation may take a while.

---

An alternative, if you have direct access to the database server, is to delete the database altogether and have Firely Server recreate it again.

- If you run on SQL Server, see *Using SQL server* for the `AutoUpdateDatabase` feature.
- If you run on MongoDB, Firely Server will recreate the collection by default if it is not present.

Although the operation requires no further arguments, it requires a POST, since it is certainly not side-effect free.

## 7.13 Preloading resources

If you have set up Firely Server as a reference server in a testing environment, it can be useful to load it with an ‘iron test set’ of examples. You can do that with the preload feature. Usually you will want to [Reset the database](#) first.

To preload a set of resources, execute:

```
POST http(s)://<firely-server-endpoint>/administration/preload
Content-Type: application/octet-stream
Body: a zip file with resources, each resource in a separate file (xml or json).
```

Firely Server will return statuscode 200 if the operation succeeded.

If you are *not permitted* to preload resources into the database, Firely Server will return statuscode 403.

**Note:** The operation can take quite long if the zip contains many resources. E.g. when uploading the [examples-json.zip](#) from the specification, it took about a minute on MongoDB and about 7 minutes on SQL Server on a simple test server.

**Attention:** This feature is not meant for bulk uploading really large sets of resources. Firely Server currently has no special operation for bulk operations.

## 7.14 Multiple versions of FHIR

Since version 3.0.0 Firely Server can run multiple versions of FHIR side-by-side in the same server. This page explains how this works and what the consequences are.

### 7.14.1 Requests

The FHIR Specification explains the mimetype parameter that distinguishes one FHIR version from another in the paragraph on the [FHIR Version parameter](#). Firely Server uses this approach to let you choose the model for your request. Below are examples on how to use the fhirVersion parameter and how it influences the behaviour of Firely Server. Accepted values for the parameter are:

- fhirVersion=3.0, for FHIR STU3
- fhirVersion=4.0, for FHIR R4

You can add the fhirVersion to the Accept and/or the Content-Type header. If you specify it on both, the fhirVersion parameters have to be the same.

**Note:** The fhirVersion parameter is also part of the Content-Type header of the response by Firely Server. Settings can control this, see [Running different versions on different endpoints](#) below.

The examples below explain the behaviour with STU3, but if you replace fhirVersion with 4.0, it works exactly the same on R4.

---

**Note:** If you do not specify a `fhirVersion` parameter, Firely Server will use `fhirVersion=3.0` (STU3) as a default. This way the behaviour is compatible with previous versions of Firely Server. If you like, you can change the Default in [Information model](#)

---

---

**Note:** If you use both an `Accept` header and a `Content-Type` header, the `fhirVersion` parameter for both must be the same. So this would be *invalid*:

```
POST <base>/Patient
Accept=application/fhir+json; fhirVersion=3.0
Content-Type=application/fhir+json; fhirVersion=4.0
```

---

Search for all Patients in STU3. In Firely Server this means Patient resources that were also stored as STU3. There is no automatic conversion of resources that were stored as R4 to the STU3 format (or vice versa).

```
GET <base>/Patient
Accept=application/fhir+json; fhirVersion=3.0
```

---

Search for Patients with the name 'Fred' in STU3. The search parameters used in the query must be valid in STU3.

```
GET <base>/Patient?name=Fred
Accept=application/fhir+json; fhirVersion=3.0
```

---

Create a Patient resource in STU3. This will only be retrievable when accessed with STU3:

```
POST <base>/Patient
Content-Type=application/fhir+json; fhirVersion=3.0
Accept=application/fhir+json; fhirVersion=3.0

{<valid Patient JSON body>}
```

---

Update a Patient resource in STU3.:

```
PUT <base>/Patient/123
Content-Type=application/fhir+json; fhirVersion=3.0
Accept=application/fhir+json; fhirVersion=3.0

{<valid Patient JSON body with id: 123>}
```

---

1. If no resource with this id existed before: it will be created with this id. (This was already always the behaviour of Firely Server.)
2. If a resource with this id existed before, in STU3: update it.
3. If a resource with this id already exists in R4: you will get an error with an `OperationOutcome` saying that a resource with this id already exists with a different `informationmodel`.

---

**Note:** Id's still have to be unique within a `resourcetype`, regardless of the FHIR version.

---

Delete a Patient resource.:

```
DELETE <base>/Patient/123
Accept=application/fhir+json; fhirVersion=3.0
```

This will delete Patient/123, regardless of its FHIR version. The Accept header is needed for Firely Server to know how to format an OperationOutcome if there is an error.

## 7.14.2 Conformance resources

Conformance resources like StructureDefinition and SearchParameter are registered *per FHIR version*. This implies:

1. Conformance resources will be imported during *Import of Conformance Resources* for both STU3 and R4. To avoid id clashes (see note above), the id's in R4 are appended with '-Fhir4.0'

1. So the StructureDefinition for Patient will be available for STU3 and R4 respectively like this:

```
GET <base>/StructureDefinition/Patient
Accept=application/fhir+json; fhirVersion=3.0

GET <base>/StructureDefinition/Patient-Fhir4.0
Accept=application/fhir+json; fhirVersion=4.0
```

2. If you add a StructureDefinition or SearchParameter via the Administration API, you can decide for yourself whether to append the FHIR version to the id or not. Just note that you cannot use the same id for different FHIR versions.
3. Depending on the fhirVersion parameter Firely Server evaluates whether a resourcetype or searchparameter is valid in that FHIR version. E.g. 'VerificationResult' is only valid in R4, but 'DataElement' is only valid in R3.
4. For validation, the StructureDefinitions and terminology resources needed are only searched for in the FHIR version of the resource that is being validated.
5. When you *Manage Conformance Resources with the Administration API*, a StructureDefinition can only be posted to the Administration API in the context of a FHIR Version that matches the StructureDefinition.fhirVersion. So this works:

```
POST <base>/administration/StructureDefinition
Accept=application/fhir+json; fhirVersion=4.0
Content-Type=application/fhir+json; fhirVersion=4.0

{
  "resourcetype": "StructureDefinition"
  ...
  "fhirVersion": "4.0.0" //Note the FHIR version matching the Content-Type
}
```

But it would not work if "fhirVersion"="3.0.1"

6. If you *Load Conformance Resources on demand*, this will be done for all the importfiles described above, regardless of the fhirVersion in the Accept header.

### 7.14.3 Running a single version

To use only a single version you set the Default information model in *Information model* to the version you want to use. In addition, you can exclude the namespace of the version you don't need (Vonk.Fhir.R3 or Vonk.Fhir.R4) from the *PipelineOptions* to disable its use. If you exclude a namespace, make sure to exclude it from all branches.

### 7.14.4 Running different versions on different endpoints

To assign endpoints to different versions, create a mapping in *Information model*. Use the Mode switch to select either a path or a subdomain mapping, assigning your endpoints in the Map array. Mapped endpoints will only accept the version you have specified. The web service root ('/' and '/administration/') will still accept all supported versions.

Assigning an endpoint to a FHIR version is exactly equivalent to adding that particular fhirVersion MIME parameter to every single request sent to that endpoint. So using these settings:

```
"InformationModel": {
  "Default": "Fhir4.0",
  "IncludeFhirVersion": ["Fhir4.0", "Fhir5.0"],
  "Mapping": {
    "Mode": "Path",
    "Map": {
      "/R3": "Fhir3.0",
      "/R4": "Fhir4.0"
    }
  }
}
```

The call

```
GET http://myserver.org/Patient
Accept=application/fhir+json; fhirVersion=3.0
```

is equivalent to

```
GET http://myserver.org/R3/Patient
```

and the call

```
GET http://myserver.org/Patient (defaults to R4)
```

is equivalent to

```
GET http://myserver.org/R4/Patient
```

and the administration call

```
GET http://myserver.org/administration/StructureDefinition (defaults to R4)
```

is equivalent to

```
GET http://myserver.org/administration/R4/StructureDefinition (/R4 is a postfix to '/
↪administration')
```

As you can see, on a mapped endpoint it is never necessary to use a FHIR `_format` parameter or a `fhirVersion` MIME parameter in a `Content-Type` or `Accept` header.



## Response Content-Type

The setting `IncludeFhirVersion` is used for the Content-Type of the response from Firely Server. Some clients cannot handle a parameter on the mimetype, and the `fhirVersion` parameter was originally not part of FHIR STU3. Therefore this settings allows you to specify for which FHIR versions this parameter should be included in the Content-Type header. By default we set it to FHIR R4 and R5, as for STU3 the `fhirVersion` may be unexpected for clients.

### 7.14.5 Support for R5 (experimental!)

By default the binaries for supporting R5 are included in the Firely Server distribution (since Firely Server (Vonk) 3.3.0). By default these binaries are not loaded. See the `PipelineOptions` in `appsettings.default`, where `Vonk.Fhir.R5` is commented out. Re-enable these in your `appsettings.instance`.

The creation of `AuditEvents` is currently not supported for R5. Please disable `Vonk.Plugin.Audit` in the root endpoint of the `PipelineOptions`.

Note that there is not yet an `errata_Fhir5.0.zip` and Firely Server will complain about that in the log. You can ignore that message.

## 7.15 \$erase resources

### 7.15.1 Description

When Firely Server receives a DELETE request for a resource, it marks it as deleted in the database which makes it hidden from search results. However, the data is still present in the database. This approach is known as *soft deletion*. This comes in handy in scenarios when you want to recover accidentally deleted data. However, there are also scenarios when you *actually* want the data to be erased from the database. For that purpose, Firely Server provides the `$erase` operation.

The `$erase` operation can be executed on a resource instance level and a resource version level.

### Examples

Use the following request to erase all versions (including the historical versions) of the `Patient/example` resource from the database.

```
POST <base-url>/Patient/example/$erase
```

Use the following request to erase the specified version `xyz` and all the older versions (based on `meta.lastUpdated`) of the `Patient/example` resource from the database.

```
POST <base-url>/Patient/example/_history/xyz/$erase
```

## 7.15.2 Appsettings

To enable the \$erase operation you will first have to make sure the plugin `Vonk.Plugin.EraseOperation.EraseOperationConfiguration` is added to the `PipelineOptions` in the appsettings.

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        ...
        "Vonk.Plugin.EraseOperation.EraseOperationConfiguration"
      ],
      "Exclude": [
        ...
      ]
    }, ...etc...
  ]
}
```

## 7.15.3 AuditEvents

- It is not allowed to erase `AuditEvents`
- `AuditEvents` for the \$erase operation will contain the list of deleted items

## 7.15.4 License

The \$erase operation is part of the core Firely Server functionality. However, to use it, you may need to request an updated license from Firely. You can use your current license file if it contains `http://fire.ly/vonk/plugins/erase`.

# 7.16 Liveness and readiness

## 7.16.1 Description

It can be useful to check whether Firely Server is still up and running, and ready to handle requests. Either just for notification, or for automatic failover. A prominent use case is to start dependent services only after Firely Server is up and running, e.g. in a docker-compose or in a Helm chart.

Firely Server provides two endpoints, each in the form of a FHIR custom operation, for different purposes:

- GET <base>/\$liveness
- GET <base>/\$readiness

These align - intentionally - with the use of liveness and readiness probes in Kubernetes.

The major difference is in the ability to handle requests. Some operations on Firely Server can trigger a long running process during which the server cannot reliably handle requests, see the *Long running tasks plugin*. If this is the case, the \$liveness operation will return a 200 OK status regardless, indicating the server is up and should not be restarted (that would just delay the long running process). The \$readiness operation however, would return 423 Locked in this case. If no long running processes are active, both operations will have the same output, namely 200 OK.

If you have assigned different endpoints to different FHIR versions (see [here](#)), you can also invoke it on each FHIR version. The result is always the same for all versions that are configured in the server. E.g:

```
GET <base-url>/$liveness
GET <base-url>/STU3/$liveness
GET <base-url>/R4/$liveness
```

### 7.16.2 Results

The \$liveness operation may return one of these http status codes:

1. 200 OK: Firely Server is up and running (but may still be blocked by a long running process).
2. 402 Payment Required: The license is expired or otherwise invalid.
3. 500 or higher: An unexpected error happened, the server is not running or not reachable (in the latter case the error originates from a component in front of Firely Server).

The \$readiness operation may return one of these http status codes:

1. 200 OK: Firely Server is up and running and ready to process requests.
2. 423 Locked: Firely Server is busy with a long running operation and cannot process requests. This could among others be a *database migration* or an *import of conformance resources*. The response will have an `OperationOutcome` with additional details.
3. 402 Payment Required: The license is expired or otherwise invalid.
4. 500 or higher: An unexpected error happened, the server is not running or not reachable (in the latter case the error originates from a component in front of Firely Server).

### 7.16.3 Configuration

Both operations should be configured in the pipeline, see the plugin reference for *\$liveness* and *\$readiness*. In the default settings this is the case. Both plugins have no further configuration in the *appsettings*.



## **SECURITY**

This section lists some general recommendations and steps to configure SMART on FHIR on Firely Server, to help secure access to the data in your server.

### **8.1 Access control and SMART on FHIR**

#### **Contents**

- *Concepts*
- *Identification and Authentication*
- *Authorization*
- *Access Control Engine*
- *Custom Authentication*
- *Other forms of Authorization*
- *Configuration*
- *Compartments*
- *Tokens*
- *Firely Auth*
- *Azure Active Directory*
- *Access Control Decisions*
- *Testing*

#### **8.1.1 Concepts**

This explanation of access control and SMART on FHIR in Firely Server requires basic understanding of the *architecture* of Firely Server, so you know what is meant by middleware components and repository interfaces. It also presumes general knowledge about authentication and OAuth2.

Access control generally consists of the following parts, which will be addressed one by one:

- Identification: Who are you? – usually a user name, login, or some identifier.

- Authentication: Prove your identification – usually with a password, a certificate or some other (combination of) secret(s) owned by you.
- Authorization: What are you allowed to read or change based on your identification?
- Access Control Engine: Enforce the authorization in the context of a specific request.

Note that Firely offers *Firely Auth* as an OAuth2 provider that is optimized to use in the context of SMART on FHIR.

### 8.1.2 Identification and Authentication

Firely Server does not authenticate users itself. It is meant to be used in an [OAuth2](#) environment in which an [OAuth2 provider](#) is responsible for the identification and authentication of users. Typically, a user first enters a Web Application, e.g. a patient portal, or a mobile app. That application interactively redirects the user to the OAuth2 provider - which is the authorization server and may or may not handle authentication as well - to authenticate, and receives an OAuth2 token back. Then, the application can do an http request to Firely Server to send or receive resource(s), and provide the OAuth2 token in the http Authentication header, thereby acting on behalf of the user. Firely Server can then read the OAuth2 token and validate it with the OAuth2 authorization server. This functionality is not FHIR specific.

### 8.1.3 Authorization

Authorization in Firely Server by default is based on [SMART on FHIR](#) and more specifically the [Scopes and Launch Context](#) defined by it. The SMART specification is released in two different version as of the date of publication: [SMART v1](#) and [SMART v2](#). Both versions define a syntax for expressing “scope”-claims within an access token, as well as providing for specialized claims for conveying context information (e.g. patient launch context claims).

These are examples of scopes and launch contexts that are recognized by Firely Server (SMART v1):

- scope=user/Observation.r: the user is allowed to read Observation resources
- scope=user/Encounter.crud: the user is allowed to read, create, update and delete Encounter resources
- scope=user/\*.rs: the user is allowed to read and search any type of resource
- scope=user/\*.cud: the user is allowed to create, update and delete any type of resource
- scope=[array of individual scopes]
- patient=123: the user is allowed access to resources in the compartment of patient 123 – see [Compartments](#).

All scopes using SMART v1 can also be expressed in SMART v2:

- scope=user/Observation.r: the user is allowed to read Observation resources
- scope=user/Encounter.cu: the user is allowed to write (create and update) Encounter resources
- scope=user/\*.r: the user is allowed to read any type of resource
- scope=user/\*.cu: the user is allowed to write (create and update) any type of resource
- scope=[array of individual scopes]
- patient=123: the user is allowed access to resources in the compartment of patient 123 – see [Compartments](#).

See section [Tokens](#) for detailed requirements regarding the structure of the access token in order to enable Firely Server to enforce these scopes.

Firely Server fully supports the syntax of SMART v1: ( 'patient' | 'user' ) '/' ( fhir-resource | '\*' ) '.' ( 'read' | 'write' | '\*' )

Additionally, the syntax of SMART v2 scopes is fully supported: ( 'patient' | 'user' | 'system' ) '/' ( fhir-resource | '\*' ) '.' ( 'c' | 'r' | 'u' | 'd' | 's' | '\*' ) ? param = value

All search capabilities supported by Firely Server can also be evaluated as part of the access scope using SMART v2. Chaining and Reverse Chaining is explicitly supported here:

- `scope=user/Observation.r?category=laboratory`: the user is allowed to read Observation resources with a category element containing the code “laboratory”
- `scope=user/*.rs?_tag=http://example.org/fhir/sid/codes|tag`: the user is allowed to read and search all resource containing a tag in Meta.tag with system “<http://example.org/fhir/sid/codes>” and code “tag”
- `scope=user/Observation.rs?encounter.id=Encounter/test`: the user is allowed to see all Observation resources linked to the Encounter with id “123”.

SMART on FHIR also defines scopes starting with ‘patient/’ instead of ‘user/’. In Firely Server these are evaluated differently. With a scope of ‘patient/’ you are required to also have a ‘patient=...’ launch context to know to which patient the user connects. As mentioned above, any request is scoped to the patient compartment and requests are rejected if the patient claim is not provided in the access token. Firely Server will additionally handle user-level scopes by checking the syntax of the SMART on FHIR scopes within the access token. It enforces that only allowed resources types are accessed and only allowed actions are executed.

**Attention:** Requests using a user-level scope are not limited a pre-defined context, e.g. a Patient compartment. Therefore all matching resources are returned to the client. It is highly advised to implement additional security measures using a custom plugin, e.g. by enforcing a certain Practitioner or Encounter context.

SMART on FHIR scopes are used to:

- configure a client: which scopes can it request?
- request authorization: client requests a set of scopes
- consent: user consents to the client using the requested scopes
- access token: Firely Server can read from the access token which scopes the client is granted (an intersection of the three above)

### 8.1.4 Access Control Engine

The Access Control Engine in Firely Server evaluates two types of authorization:

1. Type-Access: Is the user allowed to read or write resource(s) of a specific resourcetype?
2. Compartment: Is the data to be read or written within the current compartment (if any)?

As you may have noticed, Type-Access aligns with the concept of scopes, and Compartment aligns with the concept of launch context in SMART on FHIR.

The Firely Server `SmartContextMiddleware` component extracts the claims defined by SMART on FHIR from the OAuth2 token, and puts it into two classes that are then available for Access Control Decisions in all the interaction handlers (e.g. for read, search, create etc.)

SMART on FHIR defines launch contexts for Patient, Encounter and Location, extendible with others if needed. If a request is done with a Patient launch context, and the user is allowed to see other resource types as well, these other resource types are restricted by the [Patient CompartmentDefinition](#).

---

**Note:** To enable access to additional resources (outside the compartment), the client must request additional scopes.

---

### 8.1.5 Custom Authentication

You may build a plugin with custom middleware to provide authentication in a form that suits your needs. One example could be that you want to integrate [ASP.NET Core Identity](#) into Firely Server. Then you don't need the OAuth2 middleware, but instead can use the Identity framework to authenticate your users. See [Custom authorization plugin](#) for more details.

### 8.1.6 Other forms of Authorization

In [Access Control in Plugins and Facades](#) you can find the interfaces relevant to authorization in Firely Server. If your environment requires other authorization information than the standard SMART on FHIR claims, you can create your own implementations of these interfaces. You do this by implementing a [custom plugin](#). All the standard plugins of Firely Server can then use that implementation to enforce access control.

### 8.1.7 Configuration

You will need to add the Smart plugin to the Firely Server pipeline. See [Firely Server Plugins](#) for more information. In `appsettings[.instance].json`, locate the pipeline configuration in the `PipelineOptions` section, or copy that section from `appsettings.default.json` (see also [Changing the settings](#)):

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",
        ...
      ]
    }
  ]
}
```

Add `Vonk.Smart` (for SMART v1) or `Vonk.Plugin.SoFv2` (for SMART v2) to the list of included plugins. When you restart Firely Server, the Smart service will be added to the pipeline. An error will be thrown if both plugins are part of the pipeline. Please note that the SMART v2 plugin will allow the usage of the SMART v1 and SMART v2 syntax.

You can control the way Access Control based on SMART on FHIR behaves with the `SmartAuthorizationOptions` in the [Firely Server settings](#):

```
"SmartAuthorizationOptions": {
  "Enabled": true,
  "Filters": [
    {
      "FilterType": "Patient", //Filter on a Patient compartment if a 'patient' launch
      ↪scope is in the auth token
      "FilterArgument": "identifier=#patient#" //... for the Patient that has an
      ↪identifier matching the value of that 'patient' launch scope
    }
    //{
    // "FilterType": "Encounter", //Filter on an Encounter compartment if an 'encounter
    ↪' launch scope is in the auth token
    // "FilterArgument": "identifier=#encounter#" //... for the Encounter that has an
    ↪identifier matching the value of that 'encounter' launch scope
  ]
}
```

(continues on next page)



(continued from previous page)

```

    //},
    //{
    //  "FilterType": "RelatedPerson", //Filter on a RelatedPerson compartment if a
    ↪ 'relatedperson' launch scope is in the auth token
    //  "FilterArgument": "identifier=#relatedperson#" //... for the RelatedPerson that
    ↪ has an identifier matching the value of that 'relatedperson' launch scope
    //},
    //{
    //  "FilterType": "Practitioner", //Filter on a Practitioner compartment if a
    ↪ 'practitioner' launch scope is in the auth token
    //  "FilterArgument": "identifier=#practitioner#" //... for the Practitioner that
    ↪ has an identifier matching the value of that 'practitioner' launch scope
    //},
    //{
    //  "FilterType": "Device", //Filter on a Device compartment if a 'device' launch
    ↪ scope is in the auth token
    //  "FilterArgument": "identifier=#device#" //... for the Device that has an
    ↪ identifier matching the value of that 'device' launch scope
    //}
  ],
  "Authority": "url-to-your-identity-provider",
  // "AdditionalEndpoints": { //optional, only needed for certain identity provider setups
  //   "Issuers": ["additional-url-to-your-identity-provider"],
  //   "BaseEndpoints" : ["additional-url-to-your-identity-provider"]
  //},
  "Audience": "name-of-your-fhir-server" //Default this is empty
  // "ClaimsNamespace": "http://smarthealthit.org",
  "RequireHttpsToProvider": false, //You want this set to true (the default) in a
  ↪ production environment!
  "Protected": {
    "InstanceLevelInteractions": "read, vread, update, delete, history, conditional_
    ↪ delete, conditional_update, $validate",
    "TypeLevelInteractions": "create, search, history, conditional_create",
    "WholeSystemInteractions": "batch, transaction, history, search"
  },
  "TokenIntrospection": {
    "ClientId": "Firely Server",
    "ClientSecret": "secret"
  },
  "ShowAuthorizationPII": false,
  // "AccessTokenScopeReplace": "-",
  "SmartCapabilities": [
    "LaunchStandalone",
    "LaunchEhr",
    // "AuthorizePost",
    "ClientPublic",
    "ClientConfidentialSymmetric",
    // "ClientConfidentialAsymmetric",
    "SsoOpenidConnect",
    "ContextStandalonePatient",
    "ContextStandaloneEncounter",
    "ContextEhrPatient",

```

(continues on next page)

(continued from previous page)

```

    "ContextEhrEncounter",
    "PermissionPatient",
    "PermissionUser",
    "PermissionOffline",
    "PermissionOnline",
    "PermissionV1",
    //"PermissionV2",
    "ContextStyle",
    "ContextBanner"
  ]
}

```

- **Enabled:** With this setting you can disable ('false') the authentication and authorization altogether. When it is enabled ('true'), Firely Server will also evaluate the other settings. The default value is 'false'. This implies that authorization is disabled as if no SmartAuthorizationOptions section is present in the settings.
- **Filters:** Defines how different launch contexts are translated to search arguments. See [Compartment](#) for more background.
  - **FilterType:** Both a launch context and a CompartmentDefinition are defined by a resource type. Use FilterType to define for which launch context and related CompartmentDefinition this Filter is applicable.
  - **FilterArgument:** Translates the value of the launch context to a search argument. You can use any supported search parameter defined on FilterType. It should contain the name of the launch context enclosed in hashes (e.g. #patient#), which is substituted by the value of the claim.
- **Authority:** The base url of your identity provider, such that {{base\_url}}/.well-known/openid-configuration returns a valid configuration response ([OpenID Connect Discovery documentation](#)). At minimum, the jwks\_uri, token\_endpoint and authorization\_endpoint keys are required in addition to the keys required by the specification. See [Set up an Identity Provider](#) for more background.
- **AdditionalEndpoints:** Optional configuration setting. Add additional base authority endpoints that your identity provider also uses for operations that are listed in the .well-known document. The additional issuer setting will currently only extend the list of issuer urls that are valid within the issuer claim in the token passed to Firely Server. The token validation will be adjusted accordingly. Please note that it does not influence which issuer urls are allowed in the .well-known/openid-configuration document of the authorization server.
- **Audience:** Defines the name of this Firely Server instance as it is known to the Authorization server. Default is 'firelyserver'.
- **ClaimsNamespace:** Some authorization providers will prefix all their claims with a namespace, e.g. http://my.company.org/auth/user/\*.read. Configure the namespace here to make Firely Server interpret it correctly. It will then strip off that prefix and interpret it as just user/\*.read. By default no namespace is configured.
- **RequireHttpsToProvider:** Token exchange with an Authorization server should always happen over https. However, in a local testing scenario you may need to use http. Then you can set this to 'false'. The default value is 'true'.
- **Protected:** This setting controls which of the interactions actually require authentication. In the example values provided here, \$validate is not in the TypeLevelInteractions. This means that you can use POST [base-url]/Patient/\$validate without authorization. Since you only read Conformance resources with this interaction, this might make sense.
- **TokenIntrospection:** This setting is configurable when you use [reference tokens](#).
- **ShowAuthorizationPII:** This is a flag to indicate whether or not personally identifiable information is shown in logs.

- **AccessTokenScopeReplace:** With this optional setting you tell Firely Server which character replaces the / (forward slash) character in a SMART scope. This setting is needed in cases like working with Azure Active Directory (see details in section [Azure Active Directory](#)).
- **SmartCapabilities:** This setting can be used to configure [SMART capabilities](#). All capabilities listed here are supported by Firely Server, you can enable/disable specific capabilities based on your authorization server implementation.

---

**Note:** After properly configuring Firely Server to work with an OAuth2 authorization server, enabling SMART and configuring the SmartCapabilities for Firely Server, you are able to discover the SMART configuration metadata by retrieving `<base-url>/.well-known/smart-configuration`.

Please check section [Retrieve .well-known/smart-configuration](#) in the SMART specification for more details on how to request the metadata and how to interpret the response.

---

## 8.1.8 Compartments

In FHIR a [CompartmentDefinition](#) defines a set of resources ‘around’ a focus resource. For each type of resource that is linked to the focus resource, it defines the reference search parameters that connect the two together. The type of the focus-resource is in `CompartmentDefinition.code`, and the relations are in `CompartmentDefinition.resource`. The values for param in it can be read as a [reverse chain](#).

An example is the [Patient CompartmentDefinition](#), where a Patient resource is the focus. One of the related resource-types is Observation. Its params are subject and performer, so it is in the compartment of a specific Patient if that Patient is either the subject or the performer of the Observation.

FHIR defines CompartmentDefinitions for Patient, Encounter, RelatedPerson, Practitioner and Device. Although Firely Server is functionally not limited to these five, the specification does not allow you to define your own. Firely Server will use a CompartmentDefinition if:

- the CompartmentDefinition is known to Firely Server, see [Managing Conformance Resources](#) for options to provide them.
- the OAuth2 Token contains a claim with the same name as the `CompartmentDefinition.code` (but it must be lowercase).

So some of the launch contexts mentioned in SMART on FHIR map to CompartmentDefinitions. For example, the launch context ‘launch/patient’ and ‘launch/encounter’ map to the compartment ‘Patient’ and ‘Encounter’. Please note that launch contexts can be extended for any resource type, but not all resource types have a matching CompartmentDefinition, e.g. ‘Location’.

A CompartmentDefinition defines the relationships, but it becomes useful once you combine it with a way of specifying the actual focus resource. In SMART on FHIR, the launch context can do that, e.g. `patient=123`. As per the [SMART Scopes and Launch Context](#), the value ‘123’ is the value of the Patient.id. Together with the Patient CompartmentDefinition this defines a – what we call – Compartment in Firely Server:

- Patient with id ‘123’
- And all resources that link to that patient according to the Patient CompartmentDefinition.

There may be cases where the logical id of the focus resource is not known to the authorization server. Let’s assume it does know one of the Identifiers of a Patient. The Filters in the [Configuration](#) allow you to configure Firely Server to use the identifier search parameter as a filter instead of `_id`. The value in the configuration example does exactly that:

```
"Filters": [
  {
```

(continues on next page)

(continued from previous page)

```

    "FilterType": "Patient", //Filter on a Patient compartment if a 'patient' launch
    ↪scope is in the auth token
    "FilterArgument": "identifier=#patient#" //... for the Patient that has an
    ↪identifier matching the value of that 'patient' launch scope
  },
  ...
]

```

Please notice that it is possible that more than one Patient matches the filter. This is intended behaviour of Firely Server, and it is up to you to configure a search parameter that is guaranteed to have unique values for each Patient if you need that. You can always stick to the SMART on FHIR default of `_id` by specifying that as the filter:

```

"Filters": [
  {
    "FilterType": "Patient", //Filter on a Patient compartment if a 'patient' launch
    ↪scope is in the auth token
    "FilterArgument": "_id=#patient#" //... for the Patient that has an identifier
    ↪matching the value of that 'patient' launch scope
  },
  ...
]

```

But you can also take advantage of it and allow access only to the patients from a certain General Practitioner, of whom you happen to know the Identifier:

```

"Filters": [
  {
    "FilterType": "Patient", //Filter on a Patient compartment if a 'patient' launch
    ↪scope is in the auth token
    "FilterArgument": "general-practitioner.identifier=#patient#" //... for the Patient
    ↪that has an identifier matching the value of that 'patient' launch scope
  },
  ...
]

```

In this example the claim is still called 'patient', although it contains an Identifier of a General Practitioner. This is because the `CompartmentDefinition` is selected by matching its code to the name of the claim, regardless of the value the claim contains.

If multiple resources match the `Compartment`, that is no problem for Firely Server. You can simply configure the `Filters` according to the business rules in your organization.

## 8.1.9 Tokens

When a client application wants to access data in Firely Server on behalf of its user, it requests a token from the authorization server (configured as the `Authority` in the [Configuration](#)). The configuration of the authorization server determines which claims are *available* for a certain user, and also for the client application. The client app configuration determines which claims it *needs*. During the token request, the user is usually redirected to the authorization server, which might or might not be the authentication server as well, logs in and is then asked whether the client app is allowed to receive the requested claims. The client app cannot request any claims that are not available to that application. And it will never get any claims that are not available to the user. This flow is also explained in the [SMART App Authorization Guide](#).

The result of this flow should be a JSON Web Token (JWT) containing zero or more of the claims defined in SMART on FHIR. The claims can either be scopes or a launch context, as in the examples listed in [Authorization](#). This token is encoded as a string, and must be sent to Firely Server in the Authorization header of the request.

A valid access token for Firely Server at minimum will have:

- the `iss` claim with the base url of the OAuth server
- the `aud` the same value you've entered in `SmartAuthorizationOptions.Audience`
- the `scope` field with the scopes granted by this access token
- optionally, the `compartment` claim, if you'd like to limit this token to a certain compartment. Such a claim may be requested by using a context scope or launching a part of an EHR launch. See [Requesting context with scopes](#) for more details. For example, in case of Patient data access where the `launch/patient` scope is used, include the `patient` claim with the patient's id or identifier ([Compartments](#)) and make sure to request the `patient/<permissions>` scope permissions. Compartment claims combined with `user/<permissions>` claims are not taken into account.

**Warning:** Firely Server will not enforce any access control for resources outside of the specified compartment. Some compartment definitions do not include crucial resource types like 'Patient' or their corresponding resource type, i.e. all resources of this type regardless of any claims in the access token will be returned if requested. Please use this feature with caution! Additional custom access control is highly recommended.

## 8.1.10 Firely Auth

Firely provides an optimized OAuth2 provider that understands SMART on FHIR scopes and the FHIR resource types they apply to out of the box. Along with the SoF specific launch claims and all the various client authentication flows. This product is called Firely Auth and can be acquired as part of Firely Server. You can also evaluate it using a Firely Server evaluation license. See [Firely Auth](#) for more information.

## 8.1.11 Azure Active Directory

Azure Active Directory (v2.0) does not allow to define a scope with / (forward slash) in it, which is not compatible with the structure of a [SMART on FHIR scope](#). Therefore when you use AAD to provide SMART on FHIR scopes to Firely Server, you need to take the following steps

1. In a SMART scope, use another character (for instance -) instead of /. For example:

- `user/*.read` becomes `user-*.read`
- `user/*.write` becomes `user-*.write`
- `patient/Observation.r` becomes `patient-Observation.r`

If the used character (for instance -) is already in your SMART scope, then you can use \ (backward slash) to escape it.

- `patient/Observation.r?_id=Id-With-Dashes` becomes `patient-Observation.r?_id=Id\With\Dash`

If a \ (backward slash) is already in your SMART scope, then you can escape it with another \.

- `patient/Observation.r?_id=Id\With\BackwardSlash` becomes `patient-Observation.r?_id=Id\\With\\BackwardSlash`

2. Configure Firely Server which character is used in Step 1, then Firely Server will generate a proper [SMART on FHIR scope](#) and handle the request further. This can be configured via setting `AccessTokenScopeReplace`.

For the first step above, instead of doing it manually, you can deploy [SMART on FHIR AAD Proxy](#) to Azure, which helps you to replace / to - in a SMART scope when you request your access token. The other option would be to follow [Quickstart: Deploy Azure API for FHIR using Azure portal](#), check “SMART on FHIR proxy” box and use the proxy by following [Tutorial: Azure Active Directory SMART on FHIR proxy](#).

**Warning:** When you use the SMART on FHIR AAD Proxy, be careful with [SMART on FHIR v2 scopes](#). - is an allowed character within the access scope (see examples below). In those cases, the proxy simply replaces / with - and does not escape the original -, then Firely Server cannot figure out which - is original, which will result in a failed request.

- `patient/Observation.rs?category=http://terminology.hl7.org/CodeSystem/observation-category|laboratory`
- `Observation.rs?code=in=http://valueset.example.org/ValueSet/diabetes-codes`

### 8.1.12 Access Control Decisions

In this paragraph we will explain how Access Control Decisions are made for the various FHIR interactions. For the examples assume a Patient Compartment with identifier=123 as filter.

#### 1. Search

##### a. Direct search on compartment type

###### **Request**

`GET [base]/Patient?name=fred`

###### **Type-Access**

User must have read access to Patient, otherwise a 401 is returned.

###### **Compartment**

If a Patient Compartment is active, the Filter from it will be added to the search, e.g. `GET [base]/Patient?name=fred&identifier=123`

##### b. Search on type related to compartment

###### **Request**

`GET [base]/Observation?code=x89`

###### **Type-Access**

User must have read access to Observation, otherwise a 401 is returned.

###### **Compartment**

If a Patient Compartment is active, the links from Observation to Patient will be added to the search. In pseudo code: `GET [base]/Observation?code=x89& (subject:Patient.identifier=123 OR performer:Patient.identifier=123)`

##### c. Search on type not related to compartment

###### **Request**

`GET [base]/Organization`

###### **Type-Access**

User must have read access to Organization, otherwise a 401 is returned.

###### **Compartment**

No compartment is applicable to Organization, so no further filters are applied.

##### d. Search with include outside the compartment

**Request**

GET [base]/Patient?\_include=Patient:organization

**Type-Access**

User must have read access to Patient, otherwise a 401 is returned. If the user has read access to Organization, the \_include is evaluated. Otherwise it is ignored.

**Compartment**

Is applied as in case 1.a.

## e. Search with chaining

**Request**

GET [base]/Patient?general-practitioner.identifier=123

**Type-Access**

User must have read access to Patient, otherwise a 401 is returned. If the user has read access to Practitioner, the search argument is evaluated. Otherwise it is ignored as if the argument was not supported. If the chain has more than one link, read access is evaluated for every link in the chain.

**Compartment**

Is applied as in case 1.a.

## f. Search with chaining into the compartment

**Request**

GET [base]/Patient?link:Patient.identifier=456

**Type-Access**

User must have read access to Patient, otherwise a 401 is returned.

**Compartment**

Is applied to both Patient and link. In pseudo code: GET [base]/Patient?link:(Patient.identifier=456&Patient.identifier=123)&identifier=123  
In this case there will probably be no results.

2. Read: Is evaluated as a Search, but implicitly you only specify the \_type and \_id search parameters.
3. VRead: If a user can Read the current version of the resource, he is allowed to get the requested version as well.
4. Create

## a. Create on the compartment type

**Request**

POST [base]/Patient

**Type-Access**

User must have write access to Patient. Otherwise a 401 is returned.

**Compartment**

A Search is performed as if the new Patient were in the database, like in case 1.a. If it matches the compartment filter, the create is allowed. Otherwise a 401 is returned.

## b. Create on a type related to compartment

**Request**

POST [base]/Observation

**Type-Access**

User must have write access to Observation. Otherwise a 401 is returned. User must also have read access to Patient, in order to evaluate the Compartment.

**Compartment**

A Search is performed as if the new Observation were in the database, like in case 1.b. If it matches the compartment filter, the create is allowed. Otherwise a 401 is returned.

- c. Create on a type not related to compartment

**Request**

POST [base]/Organization

**Type-Access**

User must have write access to Organization. Otherwise a 401 is returned.

**Compartment**

Is not evaluated.

**5. Update**

- a. Update on the compartment type

**Request**

PUT [base]/Patient/123

**Type-Access**

User must have write access *and* read access to Patient, otherwise a 401 is returned.

**Compartment**

User should be allowed to Read Patient/123 and Create the Patient provided in the body. Then Update is allowed.

- b. Update on a type related to compartment

**Request**

PUT [base]/Observation/xyz

**Type-Access**

User must have write access to Observation, and read access to both Observation and Patient (the latter to evaluate the compartment)

**Compartment**

User should be allowed to Read Observation/123 and Create the Observation provided in the body. Then Update is allowed.

- 6. Delete: Allowed if the user can Read the current version of the resource, and has write access to the type of resource.
- 7. History: Allowed on the resources that the user is allowed to Read the current versions of (although it is theoretically possible that an older version would not match the compartment).

---

**Note:** A conditional create, update or delete (see the [FHIR http specification](#)), requires read permissions on the condition. Therefore, `user/*.write` will usually require additional `read` scopes.

---



### 8.1.13 Testing

Testing the access control functionality is possible on the [publicly hosted test server](#) as well as on a local instance.

On the public endpoint, there are 2 predefined users and 1 client, intended to be used from Postman:

- Client:
  - Client ID: `postman`
  - Client secret: `YXTHXspjK.2!rsz8jKQT`
- Auth URL: `https://auth.fire.ly/connect/authorize`
- Access Token URL: `https://auth.fire.ly/connect/token`
- User:
  - Username: `alice`
  - Password: `password`

You can test it locally using Firely Auth and Postman as a REST client. Please refer to [Firely Auth Introduction](#) for instructions:

You might also find it useful to enable more extensive authorization failure logging - Firely Server defaults to a secure setup and does not show what exactly went wrong during authorization. To do so, set the `ASPNETCORE_ENVIRONMENT` environment variable to `Development`.

## 8.2 Auditing

Firely Server can log access through the RESTful API for auditing purposes. It has 3 features:

1. Write requests and responses to a separate audit logfile.
2. Include user id and name from the JWT token (if present) in the audit log lines.
3. Write the audit information to AuditEvent resources in the Firely Server Data database.

These features can be enabled by including `Vonk.Plugins.Audit` in the pipeline.

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        "Vonk.Core",
        ...
        "Vonk.Plugin.Audit"
      ],
      ...
    },
    ...
  ]
}
```

See [Configure the pipeline](#) for more details on pipeline configuration.

At present, you can choose either to enable both file and database logging, or only database logging. To enable only database logging, replace `Vonk.Plugin.Audit` with `Vonk.Plugin.Audit.AuditEventConfiguration`. In addition, you can choose to log every call or only transaction batches. When you include a specific configuration class and want to enable username logging, you have to include `Vonk.Plugin.Audit.UsernameLoggingConfiguration`. Please see [Auditing](#) for the available options.

## 8.2.1 Filtering configuration

You can exclude requests from generating audit logs (both audit log file and audit event logging). This is helpful to reduce clutter in the logs. For example, you could exclude logging for an endpoint that is used for health monitoring of the server. The example below disables audit logging for all GET requests to `/Patient` and sub resources or operations.

```
"Audit": {
  "ExcludedRequests": [
    {
      "UrlPath": "/Patient",
      "Method": "GET"
    },
    {
      "UrlPath": "/Patient/*",
      "Method": "GET"
    }
  ]
},
```

The `UrlPath` property is required, but not otherwise checked (e.g. if it points to an existing resource). The wildcard (\*) can be used to expand matching in different ways, e.g.:

- `/Medication*` will match `/Medication`, `/MedicationRequest`, `/MedicationAdministration`, etc
- `/$*` will match all system level operations
- `/*/*/$validate` will match all validation operations on all resources

The `Method` property is optional. If left out, null, empty or given the value `*`, it will match all HTTP verbs. You can enter multiple verbs, delimited by the `|` symbol (e.g. `GET|POST`).

## 8.2.2 Audit log file configuration

### File

Configure where to put the audit log file and the format of its lines in a separate file named `audit.logsettings.json`. Just like the Firely Server application logging, the audit log also uses Serilog for logging audit events. The audit log settings are controlled in json configuration files called `audit.logsettings(.*).json`. The files are read in a hierarchy, exactly like the [appsettings files](#) are. Firely Server comes with default settings in `audit.logsettings.default.json`. You can adjust the way Firely Server logs its information by overriding these settings by either adding an additional file called `audit.logsettings.json` or `audit.logsettings.instance.json`, or in `audit.logsettings.default.json` directly. Alternatively you can control [Audit log settings with Environment Variables](#).

```
{
  "AuditLog": {
    "WriteTo": [
      {
        "Name": "Async",
```

(continues on next page)

(continued from previous page)

```

    "Args": {
      "configure": [
        {
          "Name": "File",
          "Args": {
            "path": "./audit/AuditLog.log",
            "rollingInterval": "Day",
            "fileSizeLimitBytes": "",
            "outputTemplate": "{Timestamp:yyyy-MM-dd HH:mm:ss.fff zzz} [
↪{Application}] [Audit] {RequestResponse} [Machine: {MachineName}] [ReqId: {RequestId}] ↪
↪[IP-Address: {Ip}] [Connection: {ConnectionId}] [UserId: {UserId}] [Username:
↪{Username}] [Path: {Path}] [Parameters: {Parameters}] [Action: {Action}] [Resource:
↪{Resource} Key:{ResourceKey}] [Search results: {SearchResultSummary}] [StatusCode:
↪{StatusCode}] {NewLine}"
          }
        }
      ]
    }
  }
}

```

The values that you can set for the File sink Args are:

- **path:** The location where the audit log file should be stored.
- **rollingInterval:** When this interval expires, the log system will start a new file. The start datetime of each interval is added to the filename. Valid values are **Infinite**, **Year**, **Month**, **Day**, **Hour**, **Minute**.
- **fileSizeLimitBytes:** Limit the size of the log file, which is 1GB by default. When it is full, the log system will start a new file.

The OutputTemplate listed here contains all the properties that can be logged:

- **Timestamp:yyyy-MM-dd HH:mm:ss.fff zzz:** When this was logged, with formatting.
- **Application:** Firely Server
- **RequestResponse:** indicates whether the audit event was a request or a response.
- **MachineName:** Name of the machine hosting the Firely Server instance. Especially useful when running multiple instances all logging to the same file.
- **RequestId:** Unique id of this request, use this to correlate request and response.
- **Ip:** IP Address of the client.
- **ConnectionId:** Use this to correlate requests from the same client.
- **UserId:** User id from the JWT token (if present).
- **Username:** User name from the JWT token (if present).
- **Path:** Request url.
- **Parameters:** The request parameters used.
- **Action:** Interaction that was requested (like `instance_read` or `type_search`).
- **Resource:** Resourcetype involved.

- **ResourceKey**: 'Key' of the resource involved (if any), consisting of the resourcetype and the id, formatted as "resourcetype/id".
- **StatusCode**: Statuscode of the response at the time of logging (by default '-1' when the request is not handled yet).

For transactions and batches, the audit plugin will write a line for the transaction/batch as a whole *and* one for every entry in the transaction/batch.

### Seq

Because we use Serilog for logging audit events, other Log sinks like *Seq* are also supported. *Seq* is a web interface to easily inspect structured logs.

For the Seq sink, you can also specify arguments. One of them is the server URL for your Seq server:

```
"WriteTo": [  
  {  
    "Name": "Seq",  
    "Args": { "serverUrl": "http://localhost:5341" }  
  }  
]
```

- Change `serverUrl` to the URL of your Seq server

### Audit Log reliability and performance considerations

The default Serilog log sink in `audit.logsettings.default.json` is a asynchronous wrapper around a File sink, which means that audit log messages are pushed to a background worker thread to be written to the log file on disk. This improves application performance as the writing to the audit log is non-blocking. The async wrapper uses a buffer to collect the messages that need to be logged. For the current Serilog Async implementation, the default memory buffer feeding the worker thread is capped to 10,000 items. If this limit is reached any further log events will be dropped until the buffer is below this limit again. To change the limit you can add `bufferSize` to the audit logsettings. See [Serilog.Sinks.Async](#) for more details.

In normal circumstances the buffer will regularly be flushed to the underlying sink. However, when the buffer limit does get reached the reliability of writing messages is compromised and some messages will get lost while the async wrapper tries to recover. If reliability of the auditing is very important, you might want to consider using a synchronous file sink instead. See the `audit.logsettings.default.json` for an example of a synchronous File sink configuration.

```
{  
  "AuditLog": {  
    "WriteTo": [  
      {  
        "Name": "File",  
        "Args": {  
          "path": "./audit/AuditLog.log",  
          "rollingInterval": "Day",  
          "fileSizeLimitBytes": "",  
          "outputTemplate": "{Timestamp:yyyy-MM-dd HH:mm:ss.fff zzz} [{Application}  
→] [Audit] {RequestResponse} [Machine: {MachineName}] [ReqId: {RequestId}] [IP-Address:  
→{Ip}] [Connection: {ConnectionId}] [UserId: {UserId}] [Username: {Username}] [Path:  
→{Path}] [Parameters: {Parameters}] [Action: {Action}] [Resource: {Resource} Key:  
→{ResourceKey}] [Search results: {SearchResultSummary}] [StatusCode: {StatusCode}]  
→{NewLine}"  
        }  
      }  
    ]  
  }  
}
```

(continues on next page)

(continued from previous page)

```
}  
  }  
] }  
}  
}
```

The downside is that writing to the audit log is blocking and Firely Server now has to wait on the log to finish before it can continue, which in turn affects performance. You will have to try and test what works best for your use case.

### 8.2.3 AuditEvent logging

There is no further configuration for AuditEvent logging. If you include it in the pipeline, it will start generating AuditEvent resources.

For transactions and batches the audit plugin will create an AuditEvent for the transaction/batch as a whole *and* one for every entry in the transaction/batch.

Firely Server does not allow you to update or delete the AuditEvent resources through the RESTful API so the Audit log cannot be tampered with. You can of course still manipulate these resources directly on the database, for instance to offload a surplus of old AuditEvent resources elsewhere. Please [Contact us](#) for details if you want to do this.



## FIRELY SERVER ADMINISTRATION API

Besides the regular FHIR endpoint, Firely Server also exposes an Administration API. The endpoint for this is:

`http(s)://<firely-server-endpoint>/administration`

### 9.1 Functions

The following functions are available in the Administration API.

- *Managing Conformance Resources*
- *Subscriptions*
- *Re-indexing for new or changed SearchParameters*
- *Reset the database*
- *Preloading resources*
- *Terminology services*

### 9.2 Configuration

You can *Configure the Administration API*, including restricting access to functions of the Administration API to specific ip addresses.

### 9.3 Database

The Administration API uses a database separately from the main ‘Firely Server Data’ database. Historically, SQL Server, MongoDB and Memory are supported as databases for the Administration API. As of Firely Server (Vonk) version 0.7.1, SQLite is advised for this, and we have made that the default configuration. See *Using SQLite* on how to configure for this.





## COMPLIANCE WITH NATIONAL IMPLEMENTATION GUIDES AND CERTIFICATION PROGRAMS

The following section will highlight against which national Implementation Guides and certification programs Firely Server has been successfully tested. Additionally, any necessary configuration to re-produce these testing results will be outlined.

### 10.1 Nictiz

- Tested Version: Firely Server has been tested against [nictiz.fhir.nl.stu3.zib2017 version 2.1.12](#)
- The conformance resources are **NOT** loaded by default in the standard SQLite Administration Database of Firely Server. Please see *Import of Conformance Resources* on how to load the profiles, extensions and ValueSets into Firely Server.

#### 10.1.1 Known Limitations

The following tickets represent outstanding issues with the above-mentioned specification:

- <https://bits.nictiz.nl/browse/MM-3218>
- <https://bits.nictiz.nl/browse/MM-3219>
- <https://bits.nictiz.nl/browse/MM-3221>
- <https://bits.nictiz.nl/browse/MM-3222>

These issues do not influence the validation of instances against the erroneous profiles. However, the profiles cannot be loaded via the administration REST API as the validation on that endpoint will reject the StructureDefinitions.

#### 10.1.2 Test Data

Official test data can be found in the Simplifier project for Nictiz STU3 Zib 2017.

## 10.2 GDPR compliance

Firely Server is a well-tested, secure HL7 FHIR® server that enables you to comply with the technical and organizational measures of the EU General Data Protection Regulation (GDPR).

On this page we will detail how you can achieve compliance for your Firely Server deployment. To ensure your organisation's specific use case, environment, and deployment are compliant, feel free to [contact us](#): we'd be happy to help. We also recommend checking <https://gdprchecklist.io/> for information. The following sections only focus on technical requirements, organizational requirements are out-of-scope for this document.

**Attention:** Firely is not a Data Processor as defined in by the GDPR, as Firely does not process any data stored in Firely Server.

Within the GDPR, chapter three defines a list of granted rights for a data subject. For some of these rights, Firely Server offers functionality to support the execution of these rights.

### 10.2.1 Right to rectification

Every resource in Firely Server can be updated using the Restful API, given that the executing user/client has sufficient rights and permissions. If the data subject requests rectification of their data, a data processor could update the resources containing incomplete or wrong information.

### 10.2.2 Right to erasure

Please note that the Restful “DELETE” operation is not GDPR compliant by default. After a DELETE, a resource is only marked in the database as deleted. This behaviour might be needed to comply with other (regulatory) obligations to retain records. In case it is allowed to permanently delete a resource, Firely Server offers the *\$erase operation* to delete a resource without retaining history.

### 10.2.3 Conditions for consent

FHIR offers the functionality of capturing Consent information using the [Consent resource](#). Firely Server offers a plugin framework based on which a Consent check can be implemented based on concrete business requirements.

### 10.2.4 Right to portability

All resource belonging to a data subject can be exported by either querying individual resources through the Restful API or by using the *Bulk Data Export option*. If the data subject is a Patient, all resources belonging to the Patient can be queried or exported using its Patient compartment.

## 10.3 HIPAA compliance

Firely Server is a well-tested, secure HL7 FHIR® server that enables you to comply with the Technical Safeguards of the HIPAA Security Rule.

On this page we will detail how you can achieve compliance for your Firely Server deployment. To ensure your organisation's specific usecase, environment, and deployment are compliant, feel free to [contact us](#): we'd be happy to help.

### 10.3.1 164.312(a)(1) Standard: Access control

Implement technical policies and procedures for electronic information systems that maintain electronic protected health information to allow access only to those persons or software programs that have been granted access rights as specified in 164.308(a)(4).

There are several ways to approach this:

1. ensure Firely Server is deployed in a secure environment where only those with correct permissions are able to access it,
2. use SMART on FHIR as a means of controlling access,
3. or add custom authentication based on a plugin.

Deploying in a secure environment (1) would mean access to Firely Server is controlled by third-party software or policy, placing this scenario outside the scope of this guide.

For scenario (2), Firely Server implements support for [Smart on FHIR](#), a sibling specification to FHIR for securely connecting third-party applications to Electronic Health Record data. See [Access control and SMART on FHIR](#) on how to configure Firely Server with it.

You may also wish to setup custom authentication (3). Given how Firely Server is based on a pipeline architecture, you can insert a plugin at the start of the pipeline to call out to your authentication service(s) prior to handling the request. See [this gist](#) as an example.

### 10.3.2 164.312(c)(1) Standard: Integrity

Implement policies and procedures to protect electronic protected health information from improper alteration or destruction.

The same solutions apply to this point as [164.312\(a\)\(1\) Standard: Access control](#) and [164.312\(b\) Standard: Audit control](#).

### 10.3.3 164.312(d) Standard: Person or entity authentication

Implement procedures to verify that a person or entity seeking access to electronic protected health information is the one claimed.

The same solutions apply to this point as [164.312\(a\)\(1\) Standard: Access control](#).

### 10.3.4 164.312(a)(2)(i) Unique user identification

Assign a unique name and/or number for identifying and tracking user identity.

The same solution applies to this point as *164.312(b) Standard: Audit control*. For Firely Server to be able to log the identity of the user, this identity must be present in or derivable from the authentication token, and it must be added to the log properties. If you use the SMART on FHIR plugin, that is automatically configured. If you want to do this from within a custom authentication plugin, feel free to contact us for details.

### 10.3.5 164.312(b) Standard: Audit control

Implement hardware, software, and/or procedural mechanisms that record and examine activity in information systems that contain or use electronic protected health information.

With the use of the *Audit Event log* plugin, Firely Server will thoroughly log every interaction as a note in a log file and/or in an AuditEvent resource. Logged information will be a trace record of all system activity: viewing, modification, deletion and creation of all Electronic Protected Health Information (ePHI).

The audit trail can track the source IP, event type, date/time, and more. If a JWT token is provided (for SMART on FHIR), the user/patient identity can be logged as well.

### 10.3.6 164.312(e)(1, 2) Standard: Transmission security

Implement technical security measures to guard against unauthorized access to electronic protected health information that is being transmitted over an electronic communications network.

Implement a mechanism to encrypt electronic protected health information whenever deemed appropriate.

Transmission security in Firely Server can be achieved by encrypting the communications with TLS/SSL. Standard industry practice is to use a reverse proxy (e.g. nginx or IIS) for this purpose. If you'd like, you can also enable secure connections in Firely Server *directly* without a proxy as well.

Firely Server is regularly updated with the latest versions of ASP.NET to ensure that the latest cryptographic algorithms are available for use.

### 10.3.7 164.312(e)(2)(ii) Encryption

Implement a mechanism to encrypt electronic protected health information whenever deemed appropriate.

The recommended way to ensure that e-PHI is encrypted as necessary is to use disk encryption, and there are several solutions for this depending on your deployment environment. If you're deploying in the cloud - see your vendors options for disk encryption, as most have options for encrypted disks already. If you're deploying locally, look into BitLocker on Windows or dm-crypt/LUKS for Linux.

Disk encryption is preferred over individual database field encryption as the latter would severely impact the search performance.

### 10.3.8 164.312(a)(2)(ii) Emergency access procedure

Establish (and implement as needed) procedures for obtaining necessary electronic protected health information during an emergency.

This depends on the solution you went with for *164.312(a)(1) Standard: Access control*.

In case you went with SMART on FHIR, add an authorization workflow that grants emergency access rights - essentially, a “super” access token. The application can then use this token with Firely Server, just like any other token.

If you went with a custom authentication scheme, add a special measure to handle this scenario.

### 10.3.9 164.312(a)(c) Implementation specification: Mechanism to authenticate electronic protected health information

Implement electronic mechanisms to corroborate that electronic protected health information has not been altered or destroyed in an unauthorized manner.

Firely Server does not allow you to delete resources through its RESTful API. Old versions of resources are retained by default. The only way to alter or destroy resources is through direct database access.

Therefore database-level safety mechanisms must ensure that information is not altered or destroyed unless it's desired.



## PROGRAMMING API REFERENCE & PLUGINS

Firely Server offers an extensive programming API which allows changing the behaviour of the server. Customizations can be implemented using a plugin infrastructure. The reference documentation lists the available plugins for configuring the pipeline, and the public programming API of Firely Server for building Plugins and Facades.

### 11.1 Firely Server Plugins

Firely Server Plugins is the means to adjust a Firely Server to your own special needs, beyond the configuration. Please have a look at the *Overview of Firely Server, Plugins and Facades* to see how plugins fit in the Firely Server. *Architecture* provides more detail on that.

#### 11.1.1 Configure the pipeline

Configuration of the pipeline in Firely Server is done with `PipelineOptions` in combination with `SupportedInteractions`. A default setup is installed with Firely Server in `appsettings.default.json`, and it looks like this:

```
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",
        "Vonk.Fhir.R4",
        "Vonk.Repository.Sql.SqlVonkConfiguration",
        "Vonk.Repository.Sqlite.SqliteVonkConfiguration",
        "Vonk.Repository.MongoDb.MongoDbVonkConfiguration",
        "Vonk.Repository.Memory.MemoryVonkConfiguration",
        "Vonk.Subscriptions",
        "Vonk.Smart",
        "Vonk.UI.Demo",
        "Vonk.Plugins.Terminology",
        "Vonk.Plugin.DocumentOperation.DocumentOperationConfiguration",
        "Vonk.Plugin.ConvertOperation.ConvertOperationConfiguration",
        "Vonk.Plugin.BinaryWrapper.BinaryWrapperConfiguration",
        "Vonk.Plugin.MappingToStructureMap.MappingToStructureMapConfiguration",
        "Vonk.Plugin.TransformOperation.TransformOperationConfiguration"
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    ],
    "Exclude": [
        "Vonk.Subscriptions.Administration"
    ]
},
{
    "Path": "/administration",
    "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",
        "Vonk.Fhir.R4",
        "Vonk.Repository.Sql.SqlAdministrationConfiguration",
        "Vonk.Repository.Sqlite.SqliteAdministrationConfiguration",
        "Vonk.Repository.MongoDb.MongoDbAdminConfiguration",
        "Vonk.Repository.Memory.MemoryAdministrationConfiguration",
        "Vonk.Subscriptions.Administration",
        "Vonk.Plugins.Terminology",
        "Vonk.Administration"
    ],
    "Exclude": [
        "Vonk.Core.Operations",
        "Vonk.Core.Licensing.LicenseRequestJobConfiguration"
    ]
}
],
},
"SupportedInteractions": {
    "InstanceLevelInteractions": "read, vread, update, delete, history, conditional_delete,
↪ conditional_update, $validate, $validate-code, $expand, $compose, $meta, $meta-add,
↪ $document",
    "TypeLevelInteractions": "create, search, history, conditional_create, compartment_
↪ type_search, $validate, $snapshot, $validate-code, $expand, $lookup, $compose,
↪ $document",
    "WholeSystemInteractions": "capabilities, batch, transaction, history, search,
↪ compartment_system_search, $validate, $convert"
},
},

```

**PluginDirectory:**

You can put plugins of your own (or third party) into this directory for Firely Server to pick them up, without polluting the Firely Server binaries directory itself. For a list of available plugins in Firely Server, see [Plugins available for Firely Server](#). The directory in the default setting of `./plugins` is not created upon install, you may do this yourself if you want to add a plugin.

**PluginDirectory.Branches:**

A web application can branch into different paths, and Firely Server has two by default:

- `/`: the root branch, where the main *FHIR RESTful API* is hosted;
- `/administration`: where the *Firely Server Administration API* is hosted.

Branches contains a subdocument for each of the defined paths:

**Path**

The path for this branch. This is the part after the base URL that Firely Server is hosted on.

**Include**



(Prefixes of) *Configuration classes* that add services and middleware to Firely Server.

### Exclude

(Prefixes of) *Configuration classes* that may not be executed. **Exclude** overrides **Include** and is useful if you want to use all but one configuration class from a namespace.

### SupportedInteractions:

A comma-separated list of all interactions Firely Server should enable on [base]/[type]/[id] (InstanceLevelInteractions), [base]/[type] (TypeLevelInteractions), and [base] (WholeSystemInteractions) levels. Firely Server will use this list to enable/disable supported interactions and reflect it in /metadata accordingly.

If you'd like to limit what operations your Firely Server supports, remove them from this list.

If you've added a custom plugin that enables a new interaction, make sure to load the plugin (see **PluginDirectory** above) and enable the interaction in this list. For example, if you've added the **Vonk.Plugin.ConvertOperation \$convert** plugin in **PipelineOptions.Branches.Include**, make sure to enable the operation **\$convert** as well:

```
"WholeSystemInteractions": "$convert, capabilities, batch, transaction, history,
↪search, compartment_system_search, $validate"
```

## 11.1.2 Configuration classes

A configuration class is a static class with two public static methods having the signature as below, that can add services to the Firely Server dependency injection system, and add middleware to the pipeline.

```
[VonkConfiguration (order: xyz)] //xyz is an integer
public static class MyVonkConfiguration
{
    public static void ConfigureServices(IServiceCollection services)
    {
        //add services here to the DI system of ASP.NET Core
    }

    public static void Configure(IApplicationBuilder builder)
    {
        //add middleware to the pipeline being built with the builder
    }
}
```

As you may have noticed, the methods resemble those in an ASP.NET Core Startup class. That is exactly where they are ultimately called from. We'll explain each of the parts in more detail.

### VonkConfiguration

This is an attribute defined by Firely Server (package Vonk.Core, namespace Vonk.Core.Pluggability). It tells Firely Server to execute the methods in this configuration class. The order property determines where in the pipeline the middleware will be added. You can see the order of the plugins in the *log* at startup.

### MyVonkConfiguration

You can give the class any name you want, it will be recognized by Firely Server through the attribute, not the classname. We do advise you to choose a name that actually describes what is configured. It is also better to have multiple smaller configuration classes than one monolith adding all your plugins, so you allow yourself to configure your plugins individually afterwards.

### ConfigureServices

The main requirements for this method are:

- It is public static;
- It has a first formal argument of type `Microsoft.Extensions.DependencyInjection.IServiceCollection`;
- It is the only method in this class matching the first two requirements.

This also means that you can give it a different name. Beyond that, you may add formal arguments for services that you need during configuration. You can only use services that are available from the ASP.NET Core hosting process, not any services you have added yourself earlier. Usual services to request are:

- `IConfiguration`
- `IHostingEnvironment`

These services will be injected automatically by Firely Server. See below for additional guidance on how to register services in the DI container.

### Configure

The main requirements for this method are:

- It is public static;
- It has a first formal argument of type `Microsoft.AspNetCore.Builder.IApplicationBuilder`;
- It is the only method in this class matching the first two requirements.

This also means that you can give it a different name. Beyond that, you may add formal arguments for services that you may need during configuration. Here you can use services that are available from the ASP.NET Core hosting process *and* any services you have added yourself earlier. For services in request scope please note that this method is not run in request scope. These services will be injected automatically by Firely Server.

We provided an *example* of this: creating your own landing page.

## Register a service in your plugin

Often you will want to use .NET Core provided services, or services from other common libraries in your Facade or plugin. Firely Server itself may or may not register the same service or interface already. There is a safe way to register a service if it is not registered already. The example below shows that for an `IMemoryCache`:

```
public static IServiceCollection ConfigureServices(this IServiceCollection
↪services)
{
    services.TryAddSingleton<IMemoryCache, MemoryCache>();
}

//using it in a constructor
public class MyPluginService{
    public MyPluginService(IMemoryCache cache){...}
}
```

However, should Firely Server itself have registered a service for the same interface already, you will get that one injected. Even safer is to make sure you get your own injected, e.g. by registering a derived class:

```

public class MyMemoryCache: MemoryCache{}

public static IServiceCollection ConfigureServices(this IServiceCollection
    ↪services)
{
    services.TryAddSingleton<MyMemoryCache>();
}

//using it in a constructor
public class MyPluginService{
    public MyPluginService(MyMemoryCache cache){...}
}

```

### 11.1.3 Detailed logging of loading plugins

If your plugin or any of the Firely Server plugins appears not to be loaded correctly, you may inspect what happens in more detail in the log. See [Log settings](#) for where you can find the log file. You can vary the log level for `Vonk.Core.Pluggability.VonkConfigurer` to hide or reveal details.

On the Information level, Firely Server will tell you which assemblies are loaded and searched for `VonkConfiguration` attributes:

```

Looking for Configuration in these assemblies:
C:\data\dd18\vonk_preview\Vonk.Administration.Api.dll
C:\data\dd18\vonk_preview\Vonk.Core.dll
C:\data\dd18\vonk_preview\Vonk.Fhir.R3.dll
C:\data\dd18\vonk_preview\Vonk.Fhir.R4.dll
C:\data\dd18\vonk_preview\Vonk.Repository.Generic.dll
C:\data\dd18\vonk_preview\Vonk.Repository.Memory.dll
C:\data\dd18\vonk_preview\Vonk.Repository.MongoDb.dll
C:\data\dd18\vonk_preview\Vonk.Repository.Sql.dll
C:\data\dd18\vonk_preview\vonk.server.dll
C:\data\dd18\vonk_preview\Vonk.Server.PrecompiledViews.dll
C:\data\dd18\vonk_preview\Vonk.Smart.dll
C:\data\dd18\vonk_preview\Vonk.Subscriptions.dll
C:\data\dd18\vonk_preview\Vonk.UI.Demo.dll
C:\data\dd18\vonk_preview\plugins\Visi.Repository.dll
C:\data\dd18\vonk_preview\plugins\Vonk.Facade.Relational.dll

```

Also on the Information level, Firely Server will show the services and middleware as it has loaded, in order. The list below is also the default pipeline as it is configured for Firely Server.

```

Configuration:
/
    FhirR3Configuration           [100] | Services: V | Pipeline: X
    FhirR4Configuration           [101] | Services: V | Pipeline: X
    MetadataConfiguration         [110] | Services: V | Pipeline: X
    LicenseConfiguration          [120] | Services: V | Pipeline: X
    SerializationConfiguration    [130] | Services: V | Pipeline: X
    RepositorySearchSupportConfiguration [140] | Services: V | Pipeline: X
    RepositoryIndexSupportConfiguration [141] | Services: V | Pipeline: X
    PluggabilityConfiguration     [150] | Services: V | Pipeline: X

```

(continues on next page)

(continued from previous page)

ViSiConfiguration	[240]	Services: V	Pipeline: X
DemoUIConfiguration	[800]	Services: V	Pipeline: V
VonkToHttpConfiguration	[1110]	Services: V	Pipeline: V
VonkFeaturesExtensions	[1120]	Services: X	Pipeline: V
FormatConfiguration	[1130]	Services: V	Pipeline: V
LongRunningConfiguration	[1170]	Services: V	Pipeline: V
VonkCompartmentsExtensions	[1210]	Services: X	Pipeline: V
SupportedInteractionConfiguration	[1220]	Services: V	Pipeline: V
UrlMappingConfiguration	[1230]	Services: V	Pipeline: V
ElementsConfiguration	[1240]	Services: V	Pipeline: V
FhirBatchConfiguration	[3110]	Services: V	Pipeline: V
FhirTransactionConfiguration	[3120]	Services: V	Pipeline: V
SubscriptionConfiguration	[3200]	Services: V	Pipeline: V
ValidationConfiguration	[4000]	Services: V	Pipeline: V
DefaultShapesConfiguration	[4110]	Services: V	Pipeline: V
CapabilityConfiguration	[4120]	Services: V	Pipeline: V
IncludeConfiguration	[4210]	Services: V	Pipeline: X
SearchConfiguration	[4220]	Services: V	Pipeline: V
ProfileFilterConfiguration	[4310]	Services: V	Pipeline: V
PrevalidationConfiguration	[4320]	Services: V	Pipeline: V
ReadConfiguration	[4410]	Services: V	Pipeline: V
CreateConfiguration	[4420]	Services: V	Pipeline: V
UpdateConfiguration	[4430]	Services: V	Pipeline: V
DeleteConfiguration	[4440]	Services: V	Pipeline: V
ConditionalCreateConfiguration	[4510]	Services: V	Pipeline: V
ConditionalUpdateConfiguration	[4520]	Services: V	Pipeline: V
ConditionalDeleteConfiguration	[4530]	Services: V	Pipeline: V
HistoryConfiguration	[4610]	Services: V	Pipeline: V
VersionReadConfiguration	[4620]	Services: V	Pipeline: V
InstanceValidationConfiguration	[4840]	Services: V	Pipeline: V
SnapshotGenerationConfiguration	[4850]	Services: V	Pipeline: V
/administration			
SqlVonkConfiguration	[220]	Services: V	Pipeline: X
SqlAdministrationConfiguration	[221]	Services: V	Pipeline: X
DatabasePluggabilityConfiguration	[300]	Services: V	Pipeline: X
VonkToHttpConfiguration	[1110]	Services: V	Pipeline: V
VonkFeaturesExtensions	[1120]	Services: X	Pipeline: V
FormatConfiguration	[1130]	Services: V	Pipeline: V
SecurityConfiguration	[1150]	Services: V	Pipeline: V
AdministrationOperationConfiguration	[1160]	Services: V	Pipeline: V
LongRunningConfiguration	[1170]	Services: V	Pipeline: V
VonkCompartmentsExtensions	[1210]	Services: X	Pipeline: V
SupportedInteractionConfiguration	[1220]	Services: V	Pipeline: V
UrlMappingConfiguration	[1230]	Services: V	Pipeline: V
ElementsConfiguration	[1240]	Services: V	Pipeline: V
DefaultShapesConfiguration	[4110]	Services: V	Pipeline: V
AdministrationSearchConfiguration	[4221]	Services: V	Pipeline: V
ValidationConfiguration	[4310]	Services: V	Pipeline: X
SubscriptionValidationConfiguration	[4330]	Services: V	Pipeline: V
ChangeInterceptionConfiguration	[4390]	Services: X	Pipeline: V
AdministrationReadConfiguration	[4411]	Services: V	Pipeline: V
AdministrationCreateConfiguration	[4421]	Services: V	Pipeline: V

(continues on next page)

(continued from previous page)

AdministrationUpdateConfiguration	[4431]	Services: V   Pipeline: V
AdministrationDeleteConfiguration	[4441]	Services: V   Pipeline: V
AdministrationImportConfiguration	[5000]	Services: V   Pipeline: V
CodeSystemLookupConfiguration	[5110]	Services: V   Pipeline: V
ValueSetValidateCodeInstanceConfiguration	[5120]	Services: V   Pipeline: V
ValueSetValidateCodeTypeConfiguration	[5130]	Services: V   Pipeline: V
ValueSetExpandInstanceConfiguration	[5140]	Services: V   Pipeline: V
ValueSetExpandTypeConfiguration	[5150]	Services: V   Pipeline: V
CodeSystemComposeInstanceConfiguration	[5160]	Services: V   Pipeline: V
CodeSystemComposeTypeConfiguration	[5170]	Services: V   Pipeline: V

It shows all the configuration classes it found, and whether a `ConfigureServices` and / or a `Configure` method was found and executed. It also displays the value of the `Order` property of the `VonkConfiguration` attribute for each configuration class. This allows you to determine an appropriate order for your own configuration class.

On the `Verbose` level, Firely Server will also tell you why each configuration class that is found is being included or excluded. An example:

```
2018-07-02 12:58:10.586 +02:00 [Firely Server] [Verbose] [Machine: XYZ] [ReqId: ] ↳
↳ Searching for configurations in assembly "Vonk.Core, Version=0.7.0.0, Culture=neutral,
↳ PublicKeyToken=null"
2018-07-02 12:58:10.625 +02:00 [Firely Server] [Verbose] [Machine: XYZ] [ReqId: ] "Vonk.
↳ Core.Serialization.SerializationConfiguration" was included on "/" because it matches
↳ the include "Vonk.Core"
2018-07-02 12:58:10.625 +02:00 [Firely Server] [Verbose] [Machine: XYZ] [ReqId: ] "Vonk.
↳ Core.Serialization.SerializationConfiguration" was not included on "/administration"
↳ because it did not match any include
```

### 11.1.4 The order of plugins

Firely Server is organized as a pipeline of components - called `Middleware`. Every request travels through all the components until one of the components provides the response to the request. After that, the response travels back through all the components, in reverse order. Components that come *after* the responding component in the pipeline are not visited at all.

So let's say you issue a FHIR read interaction, `GET <base-url>/Patient/example`. The component implementing this interaction sits in the pipeline after search but before create. So the request will visit the search middleware (that will ignore it and just pass it on) but will never visit the create middleware.

So many components implement an interaction and provide a response to that interaction. In Firely Server those are called `Handlers`. Some components may not provide responses directly, but read or alter the request on the way in. Such a component is called a `PreHandler`. Reversely, a component may read or alter the response on the way back. Such a component is called a `PostHandler`.

A plugin can configure its own component in this pipeline but as you may understand by now it makes a difference *where* in the pipeline you put that component. Especially if it is a `Pre-` or `PostHandler`. To control the position in the pipeline, Firely Server uses the concept of `Order`.

The *`VonkConfiguration` attribute* lets you define an `Order` for your component. This page explains how to choose a suitable number for that order.

### Inspect order numbers in use

When you start Firely Server, the log lists all the loaded plugins, with their order. You can see an example [here](#). Also the list of *Plugins available for Firely Server* includes the order number chosen for each of those plugins.

### Minimum order

### Registering services

The order is mainly relevant for middleware that you register in the pipeline, in the `Configure(IApplicationBuilder app)` method. Some plugins, like e.g. a Facade implementation, only need to register services for the dependency injection framework, in the `ConfigureServices(IServiceCollection services)` method. If that is the case, the order is only relevant if you need to *override* a registration done by Firely Server. There are two ways:

1. Choose an order before Firely Server's default registration. Firely Server in general uses `TryAddSingleton` or `TryAddScoped` to register an implementation of an interface. This means that if an implementation is already registered, the `TryAdd...` will not register a second implementation.

As an example: if you want to override the registration of `IReadAuthorizer`: that is registered from the *RepositorySearchSupport* plugin, with order 140. So you would choose an order lower than 140.

2. Choose a high order (e.g. > 10000) and make sure your registration overwrites any existing registration.

```
services.AddOrReplace<IReadAuthorizer, MyReadAuthorizer>(ServiceLifetime.Scoped);
```

The latter method is the least error prone and therefore recommended.

If you only need to register interfaces and/or classes defined by your plugin, the order is not relevant, so pick any number. All service registrations are done before the pipeline itself is configured.

### Registering middleware

For middleware it is more important where exactly it ends up in the pipeline. This depends mostly on what type of handler it is, see below at *Handlers and pre- and posthandlers*.

No matter what handler you have, it probably wants to act on the `IVonkContext`. Then it is important to be in the pipeline *after* the *HttpToVonk* plugin (order: 1110), since this plugin translates information from the `HttpContext` to an `IVonkContext` and adds the latter as a feature to the `HttpContext.Features` collection.

Also, you probably want to set your response on the `IVonkContext.Response` and not directly on the `HttpContext.Response`. Then, you will need the *VonkToHttp* plugin (order: 1120) to translate the `IVonkContext` back to the `HttpContext`.

So in general, the minimum order you need for your plugin will be higher than 1120.

If you want your middleware to act on all the entries in a Batch or Transaction, you need to choose an order higher than that of the *Transaction* plugin, which is 3120.

## Order collisions

If two plugins have the same order number, it is not defined in what order the plugins will be put in the pipeline. As long as those plugins act on disjoint sets of requests that may not be a problem. But it is recommended to avoid this by checking the orders already in use.

## Handlers and pre- and posthandlers

In Firely Server you can define different types of middleware:

- Handler - acts on requests of a certain type, provides the response to it and ends the pipeline.
- Prehandler - acts on requests of certain type(s), may modify the request and sends the request further down the pipeline.
- Posthandler - lets the request pass by to be handled further down the pipeline. When the response passes on the way back, it acts on requests or responses of certain type(s), and may modify the response.

This is explained in the [session on Plugins](#) from [DevDays 2018](#).

What type of middleware you want your service to be is defined by your use of one of the `*Handle...` methods from the [VonkAppBuilder extension methods](#) or the [IApplicationBuilder extension methods](#).

### Prehandler

A Prehandler needs to act *before* the actual handler will provide a response. So the order of it must be lower than any Handler that may handle the requests that this Prehandler is interested in.

So if you want a Prehandler to intercept all create interactions, you should choose an order lower than that of the [Create](#) plugin, which is 4420.

An example of this is the [Prevalidation](#) plugin. It needs to validate all resources that get handled by the Create, Update, Conditional Create and Conditional Update plugins. Of these, Create has the lowest order: 4420. So it must be below 4420. But it also needs to act on each resource in a [Batch](#) or [Transaction](#), so it must be higher than these two, which means higher than 3120. So this is why we have chosen 4320 as order for Prevalidation.

### Posthandler

A Posthandler needs to act *after* the actual handler provided a response. But due to the nature of the processing pipeline that means it must have an order *lower* than that of the handler(s) it wants to post-process. The idea is that the posthandler sits in the pipeline and lets the request pass through. Then one of the handlers provides the response and sends it back through the pipeline. It will pass through the posthandler again (now 'backwards'), and then the posthandler will do its processing.

So if you want a Posthandler to process the responses of all create interactions (e.g. for logging purposes), you should choose an order lower than that of the [Create](#) plugin.

An example for this is the [Include](#) plugin. This must act on the response of the [Search](#) plugin. So the Include has order 4210, right before Search which has 4220.



### 11.1.5 Important classes and interfaces

If you want to develop a plugin for Firely Server, there are a couple of classes that you will probably interact with. These classes are listed under *Reference - Programming API*.

### 11.1.6 Template for a plugin

**Attention:** A complete template for a Firely Server plugin can be found on [Github](#). It covers all details on how to create a custom operation and how to use Firely Server services internally.

A regular Firely Server plugin acts on the *IVonkContext* and its *IVonkRequest* and *IVonkResponse* properties.

You don't have to create ASP.NET Core middleware yourself. You just need to create a service acting on *IVonkContext*. In the configuration you can specify when the service should be called, and in which position in the pipeline it should be put. See *Interaction Handling* for details on that.

You can use the following code as a template for a plugin:

```
using Microsoft.AspNetCore.Builder;
using Microsoft.Extensions.DependencyInjection;
using Microsoft.Extensions.DependencyInjection.Extensions;
using System.Threading.Tasks;
using Vonk.Core.Context;
using Vonk.Core.Context.Features;
using Vonk.Core.Pluggability;
using Vonk.Fhir.R3;
using F = Hl7.Fhir.Model;

namespace com.mycompany.vonk.myplugin
{
    public class MyPluginService
    {
        public async Task Act(IVonkContext vonkContext)
        {
            var (request, args, response) = vonkContext.Parts();
            //do something with the request
            //write something to the response
            response.Payload = new F.Patient{Id = "pat1"}.ToIResource();
            response.HttpResult = 200;
        }
    }

    [VonkConfiguration(order: 5000)]
    public static class MyPluginConfiguration
    {
        public static IServiceCollection AddMyPluginServices(IServiceCollection services)
        {
            services.TryAddScoped<MyPluginService>();
            return services;
        }
    }
}
```

(continues on next page)



(continued from previous page)

```

    public static IApplicationBuilder UseMyPlugin(IApplicationBuilder app)
    {
        app.OnCustomInteraction(VonkInteraction.system_custom, "myOperation").
        ↪HandleAsyncWith<MyPluginService>((svc, context) => svc.Act(context));
        return app;
    }
}

```

### 11.1.7 Returning non-FHIR content from a plugin

Some plugins may need to return content that is not a FHIR Resource. You currently cannot do that through the `IVonkResponse`. But there is another way: write directly to the `HttpContext.Response`. The plugin with order 1110 makes the `IVonkContext` accessible. That means if you pick an order higher than 1110 you can read the `IVonkContext.Request` and `.Arguments`, and write directly to the `HttpContext.Response`.

If you write the response yourself, you also need to set the `StatusCode` and `Content-Type` on the `HttpContext.Response`.

---

**Note:** If you write to the `HttpContext.Response` directly, the payload of the `IVonkContext.Response` is ignored.

---

The steps to take are:

1. Configure your plugin with an order higher than 1110
2. Write the `HttpContext.Response.Body` directly
3. Set other properties of the `HttpContext.Response` (like `StatusCode`) yourself.

An example of such a plugin would look like this. Note that this is now regular ASP.NET Core Middleware, not a service like in *Template for a plugin*.

```

using Microsoft.AspNetCore.Builder;
using Microsoft.AspNetCore.Http;
using Microsoft.Extensions.DependencyInjection;
using System.Text;
using System.Threading.Tasks;
using Vonk.Core.Context;
using Vonk.Core.Context.Features;
using Vonk.Core.Pluggability;

namespace com.mycompany.vonk.myplugin
{
    public class MyPluginMiddleware
    {
        private readonly RequestDelegate _next;

        public MyPluginMiddleware(RequestDelegate next)
        {
            _next = next;
        }

        public async Task Invoke(HttpContext httpContext)

```

(continues on next page)

(continued from previous page)

```

    {
        var vonkContext = httpContext.Vonk();
        var (request, args, _) = vonkContext.Parts();
        if (VonkInteraction.system_custom.HasFlag(request.Interaction))
        {
            //write something directly to the HttpContext.Response. Now you also
            ↪have to set the Content-Type header and the Content-Length yourself.
            var message = "This is a response that is not a FHIR resource";
            string contentLength = Encoding.UTF8.GetByteCount(message).ToString();

            httpContext.Response.Headers.Add("Content-Type", "text/plain;
            ↪charset=utf-8");
            httpContext.Response.Headers.Add("Content-Length", contentLength);
            httpContext.Response.StatusCode = 200;
            await httpContext.Response.WriteAsync("This is a response that is not a
            ↪FHIR resource");
        }
        else
        {
            await _next(httpContext);
        }
    }
}

[VonkConfiguration(order: 1115)] //note the order: higher than 1110
public static class MyPluginConfiguration
{
    public static IServiceCollection AddMyPluginServices(IServiceCollection services)
    {
        //No services to register in this example, but if you create services to do
        ↪the actual work - register them here.
        return services;
    }

    public static IApplicationBuilder UseMyPlugin(IApplicationBuilder app)
    {
        app.UseMiddleware<MyPluginMiddleware>(); //You cannot use the extension
        ↪methods that allow you to filter the requests.
        return app;
    }
}
}

```

## Custom authorization plugin

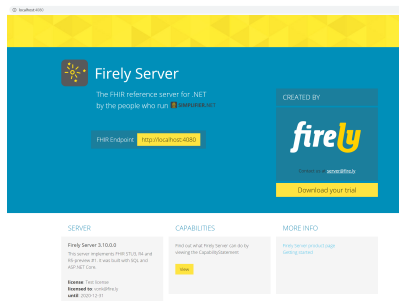
This feature can also be used to implement custom authorization. You can find a template for that in this [gist](#).

If you just return a statuscode, you could use the regular `IVonkContext.Response`. If you want to return e.g. a custom json object, you should use the method described above. Do not forget to set the Content-Type (to `application/json` for a custom json object).

For more information about access control in plugins and facades, see *Access Control in Plugins and Facades*.

### 11.1.8 Firely Server Plugin example - Create a new landing page

As a minimal example of how to use Firely Server Plugins we will show you how to create a library with your own landing page, and use it to replace the landing page that is provided by Firely Server. The landing page is the webpage you see when you access Firely Server's endpoint with a browser. By default it looks like this:



## Create a new ASP.NET Core web application

In Visual Studio create a new project of type ASP .NET Core Web Application:

## Create a new project

Recent project templates

- ASP.NET Core Web Application C#
- xUnit Test Project (.NET Core) C#
- MSTest Test Project (.NET Core) C#
- Class Library (.NET Standard) C#
- Console App (.NET Core) C#


asp.net core


Clear all


C#


All platforms


All project types

 **ASP.NET Core Web Application**  
Project templates for creating ASP.NET Core web apps and web APIs for Windows, Linux and macOS using .NET Core or .NET Framework. Create web apps with Razor Pages, MVC, or Single Page Apps (SPA) using Angular, React, or React + Redux.  
C# Linux macOS Windows Cloud Service Web

 **Console App (.NET Core)**  
A project for creating a command-line application that can run on .NET Core on Windows, Linux and MacOS.  
C# Linux macOS Windows Console

 **MSTest Test Project (.NET Core)**  
A project that contains MSTest unit tests that can run on .NET Core on Windows, Linux and MacOS.  
C# Linux macOS Windows Test

 **NUnit Test Project (.NET Core)**  
A project that contains NUnit tests that can run on .NET Core on Windows, Linux and MacOS.  
C# Linux macOS Windows Desktop Test Web

 **WPF App (.NET Core)**  
Windows Presentation Foundation client application

Back

Next

Press OK to continue.

## Configure your new project

ASP.NET Core Web Application C# Linux macOS Windows Cloud Service Web


Project name

WebApplication2

Location

C:\source\repos

...

Solution name 

WebApplication2

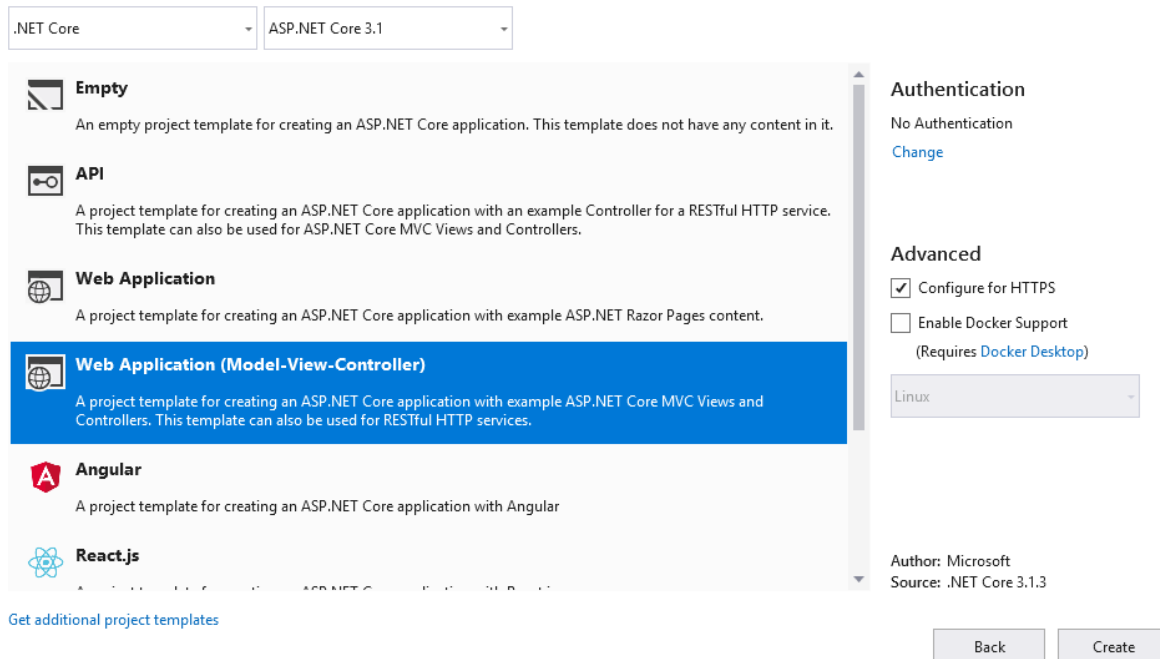
☐ Place solution and project in the same directory

Back

Create

Choose a name for your project and solution. Click Create to continue.

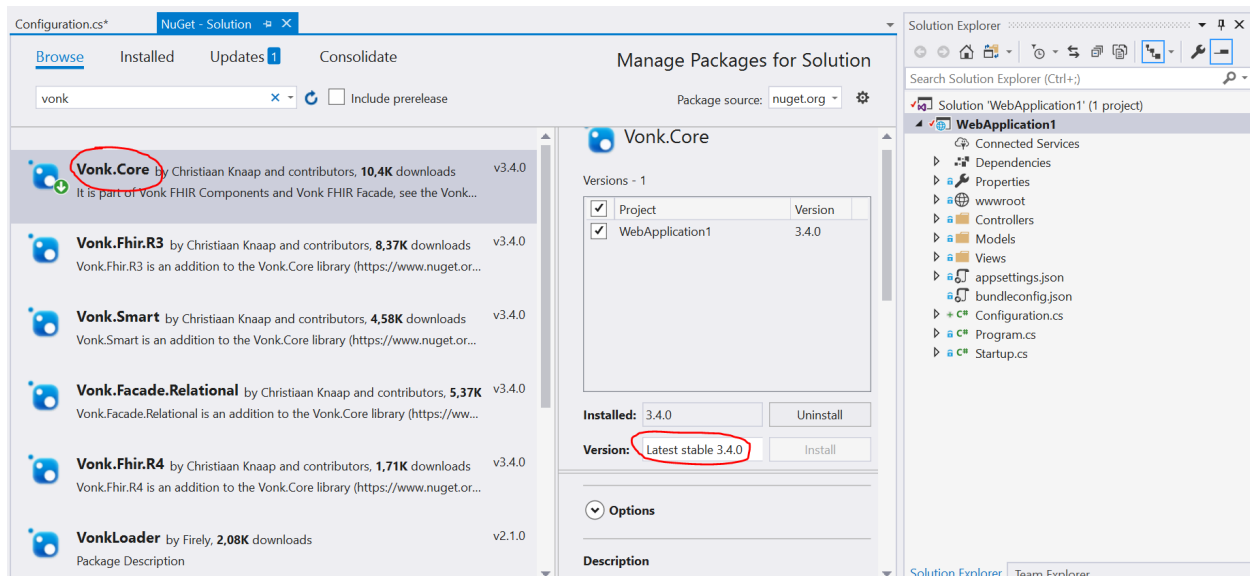
## Create a new ASP.NET Core web application



Choose ASP.NET Core 3.1 and select Web Application (Model-View-Controller). Press OK.

## Add Firely Server Package

Add Vonk.Core via the Nuget Package Manager:



This will give you access to all the core components of Firely Server, including the Vonk.Core.Pluggability.

VonkConfiguration attribute.

### Adjust project file

Add wwwroot and Views as an Embedded resource in the project file (that is necessary for Firely Server to pick them up from a library dll). To edit the project file, right click on the project file and select Edit <projectname>.cproj:

```
<ItemGroup>
  <EmbeddedResource Include="wwwroot\**\*" />
  <EmbeddedResource Include="Views\**" />
</ItemGroup>
```

The project file will look like this:

```
<Project Sdk="Microsoft.NET.Sdk.Web">

  <PropertyGroup>
    <TargetFramework>netcoreapp3.1</TargetFramework>
  </PropertyGroup>

  <ItemGroup>
    <PackageReference Include="Vonk.Core" Version="3.4.0" />
  </ItemGroup>

  <ItemGroup>
    <EmbeddedResource Include="wwwroot\**\*" />
    <EmbeddedResource Include="Views\**" />
  </ItemGroup>

</Project>
```

Save the project file.

### Create the configuration class

Next, add a new file for the configuration class, as described in *Configuration classes*. Annotate it with [VonkConfiguration(order: 802)]. In the example below the class is named UIConfiguration. Then add the static methods as prescribed:

```
public static IServiceCollection AddUIServices(IServiceCollection services)
{
    var thisAssembly = typeof(UIConfiguration).GetTypeInfo().Assembly;
    services
        .AddMvc(option => option.EnableEndpointRouting = false)
        .AddRazorRuntimeCompilation()
        .AddApplicationPart(thisAssembly)
        .AddControllersAsServices();

    var embeddedFileProvider = new EmbeddedFileProvider(
        thisAssembly,
        thisAssembly.GetName().Name
    );
}
```

(continues on next page)

(continued from previous page)

```
services.Configure<MvcRazorRuntimeCompilationOptions>(options =>
{
    options.FileProviders.Clear();
    options.FileProviders.Add(embeddedFileProvider);
});
return services;
}
```

```
public static IApplicationBuilder UseUI(IApplicationBuilder app)
{
    var thisAssembly = typeof(UIConfiguration).GetTypeInfo().Assembly;
    var embeddedStaticFileProvider = new EmbeddedFileProvider(
        thisAssembly,
        thisAssembly.GetName().Name + ".wwwroot"
    );

    app.UseStaticFiles(new StaticFileOptions() { FileProvider = ↵
↵embeddedStaticFileProvider });

    return app.MapWhen(ctx => ctx.IsBrowserRequest(), ab => ab.UseMvcWithDefaultRoute());
}
```

The source file will then look like this:

```

namespace WebApplication2
{
    [VonkConfiguration(order: 802)]
    2 references
    public static class UIConfiguration
    {
        0 references
        public static IServiceCollection AddUIServices(IServiceCollection services)
        {
            var thisAssembly = typeof(UIConfiguration).GetTypeInfo().Assembly;
            services
                .AddMvc(option => option.EnableEndpointRouting = false)
                .AddRazorRuntimeCompilation()
                .AddApplicationPart(thisAssembly)
                .AddControllersAsServices();

            var embeddedFileProvider = new EmbeddedFileProvider(
                thisAssembly,
                thisAssembly.GetName().Name
            );

            services.Configure<MvcRazorRuntimeCompilationOptions>(options =>
            {
                options.FileProviders.Clear();
                options.FileProviders.Add(embeddedFileProvider);
            });
            return services;
        }

        0 references
        public static IApplicationBuilder UseUI(IApplicationBuilder app)
        {
            var thisAssembly = typeof(UIConfiguration).GetTypeInfo().Assembly;
            var embeddedStaticFileProvider = new EmbeddedFileProvider(
                thisAssembly,
                thisAssembly.GetName().Name + ".wwwroot"
            );

            app.UseStaticFiles(new StaticFileOptions { FileProvider = embeddedStaticFileProvider });

            return app.MapWhen(ctx => ctx.IsBrowserRequest(), ab => ab.UseMvcWithDefaultRoute());
        }
    }
}

```

## Deploy and Configure

Build this project in Release mode and copy the produced dll (located in <src>\bin\Release\netcoreapp3.1) to the plugin directory of Firely Server, as configured in the *PipelineOptions:PluginDirectory*.

Go to the *Firely Server settings* of Firely Server, and replace the namespace of the landingpage (Vonk.UI.Demo) in the include of the PipelineOptions:

```

"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        "Vonk.Core",
        "Vonk.Fhir.R3",

```

(continues on next page)



(continued from previous page)

```

        "Vonk.Fhir.R4",
        //"Vonk.Fhir.R5"
        "Vonk.Repository.SqlVonkConfiguration",
        "Vonk.Repository.SqliteVonkConfiguration",
        "Vonk.Repository.MongoDbVonkConfiguration",
        "Vonk.Repository.MemoryVonkConfiguration",
        "Vonk.Subscriptions",
        "Vonk.Smart",
        "WebApplication2" //This is the adjustment you make.
        "Vonk.Plugin.DocumentOperation.DocumentOperationConfiguration",
        "Vonk.Plugin.ConvertOperation.ConvertOperationConfiguration",
        "Vonk.Plugin.BinaryWrapper",
        "Vonk.Plugin.MappingToStructureMap.MappingToStructureMapConfiguration",
        "Vonk.Plugin.TransformOperation.TransformOperationConfiguration",
        "Vonk.Plugin.Audit"
    ],
    "Exclude": [
        "Vonk.Subscriptions.Administration"
    ]
},
{
    "Path": "/administration",
    "Include": [
        "Vonk.Fhir.R3",
        "Vonk.Fhir.R4",
        //"Vonk.Fhir.R5"
        "Vonk.Repository.SqlAdministrationVonkConfiguration",
        "Vonk.Repository.SqliteAdministrationVonkConfiguration",
        "Vonk.Repository.MongoDbAdministrationVonkConfiguration",
        "Vonk.Repository.MemoryAdministrationVonkConfiguration",
        "Vonk.Subscriptions.Administration",
        "Vonk.Plugins.Terminology",
        "Vonk.Plugin.Audit",
        "Vonk.Administration"
    ],
    "Exclude": [
        "Vonk.Core.Operations",
        "Vonk.Core.Licensing.LicenseRequestJobConfiguration"
    ]
}
]
}
}

```

### Run and admire

Now run Firely Server from the commandline or Powershell window with

```
> dotnet .\Vonk.Server.dll
```

Open a browser and visit the homepage of Firely Server (<http://localhost:4080>) to admire your own landingpage.

### 11.1.9 Firely Server Plugin example - \$document operation

As a more elaborate example we implemented the \$document operation from the FHIR Specification as a plugin.

You can find the code on [GitHub](#).

### 11.1.10 BinaryWrapper plugin

#### Description

Enables you to send binary content to Firely Server and have it stored as a Binary resource, as well as the reverse: get a Binary resource and have it returned in its original binary format. The contents are Base64 encoded and stored inside the resource in the Firely Server database. Therefore this plugin is only suitable for small binary files.

Sending binary content example request:

```
POST <base-url>/Binary
Content-Type = application/pdf
Accept = application/fhir+json; fhirVersion=4.0
Body: enclose a file with the actual contents
```

- This can also be done with a PUT: PUT <base-url>/Binary/example
- The Content-Type must be the mediatype of the actual contents. It will only be accepted by Firely Server if it is one of the mediatypes listed in the RestrictToMimeTypes below.
- The Accept header can be either a fhir mediatype, with any of the supported FHIR versions. You could also set it equal to the Content-Type in which case you will be returned the same contents again.

Getting binary content example request:

```
GET <base-url>/Binary/example
Accept = application/pdf; fhirVersion=4.0
```

- The Accept header should match the mediatype of the actual contents. If you don't know the mediatype, you could request the binary resource in FHIR format first and examine its contentType element.

## Configuration

```
"Vonk.Plugin.BinaryWrapper":{
  "RestrictToMimeTypes": ["application/pdf", "text/plain", "image/png", "image/jpeg"]
},
"SizeLimits": {
  "MaxResourceSize": "1MiB", // b/kB/KiB/Mb/MiB, no spaces
},
"PipelineOptions": {
  "Branches": [
    {
      "Path": "/",
      "Include": [
        ...
        "Vonk.Plugin.BinaryWrapper"
      ]
    },
    ...
  ]
}
```

- `RestrictToMimeTypes` protects Firely Server from arbitrary content.
- This plugin honours the Firely Server setting for maximum resource size. This protects Firely Server from binary contents that are too large to store in the database.
- The namespace `Vonk.Plugin.BinaryWrapper` configures both encoding and decoding of binary contents. You can configure them separately as well:

```
"PipelineOptions": {
  "Branches": [
    {
      "Path": "/",
      "Include": [
        ...
        "Vonk.Plugin.BinaryWrapper.BinaryEncodeConfiguration",
        "Vonk.Plugin.BinaryWrapper.BinaryDecodeConfiguration"
      ]
    },
    ...
  ]
}
```

## Release notes

### Version 0.3.0

- Built against Firely Server (Vonk) 3.2.0
- Compatible with Firely Server (Vonk) 3.2.0, 3.2.1, 3.3.0
- Introduces the decoding of Binary resources, so you can GET a Binary resource in its original binary format.

### Version 0.2.0

- Build against Firely Server (Vonk) 3.0.0
- Compatible with Firely Server (Vonk) 3.0.0
- Functionally equivalent to version 0.1.0

### Version 0.1.0

- Build against Firely Server (Vonk) 2.1.0
- Compatible with Firely Server (Vonk) 2.1.0
- Introduces the encoding of Binary resources, so you can POST binary contents and have it stored as a Binary resource.

## 11.1.11 Convert plugin

### Description

Enables you to convert between json and xml representation of a resource using the \$convert operation. It can not convert between different FHIR versions.

Convert example request:

```
POST <base-url>/$convert
Content-Type = application/fhir+json; fhirVersion=4.0
Accept = application/fhir+xml; fhirVersion=4.0
Body: a resource in JSON format
```

- This can also be done reversely, with a body in XML format

### Configuration

```
"PipelineOptions": {
  "Branches": [
    {
      "Path": "/",
      "Include": [
        ...
        "Vonk.Plugin.ConvertOperation"
      ]
    },
    ...
  ]
}
```

## Release notes

### Version 0.2.0

- Build against Firely Server (Vonk) 3.2.0
- Compatible with Firely Server (Vonk) 3.2.0, 3.3.0
- Functionally equivalent to version 0.1.0

### Version 0.1.0

- Build against Firely Server (Vonk) 3.0.0
- Compatible with Firely Server (Vonk) 3.0.0, 3.1.0
- Introduces the implementation of \$convert for conversion between json and xml.

## 11.2 Firely Server Facade

Firely Server Facade is a means to use the Firely Server implementation of the FHIR RESTful API on top of an existing repository. This repository may be a relational database, a nosql database or even another web api.

This chapter details the two options for setting up a Firely Server Facade, and provides you with an exercise to get some hands-on experience with a Facade solution.

### 11.2.1 Facade setup configuration

To set up a Firely Server Facade, you will need to create a library with your own implementation of the interfaces for reading and/or writing FHIR resources and provide that as a plugin to Firely Server.

#### Provide a plugin to Firely Server

This leverages the capabilities of *Firely Server Plugins*. With this setup you:

- create a new ASP.NET Core library
- include Firely Server NuGet packages
- implement your own repository backend to interface with your data store (can be SQL server or any other medium)
- configure the PipelineOptions to use your library instead of Firely Server's own repository implementation
- configure the PipelineOptions to limit the plugins to those that are supported by your repository implementation.

The benefit of using this approach is that you automatically get to use all of Firely Server's configuration, logging, Application Insights integration, the *Administration API*, etc. described in the other sections of this documentation.

The *exercise* below uses this setup.

---

**Note:** Although we take care to try and avoid breaking changes, please be prepared to retest and update your plugins when you choose to update Firely Server.

---

### 11.2.2 Exercise: Build your first Facade

The best way to get experience in developing a Firely Server Facade is by following the exercise - build your first facade. This exercise builds a facade on a simple relational database, by creating a plugin and inserting that into the Firely Server pipeline.

Using a Firely Server Facade allows you to open up legacy systems to the FHIR ecosystem, or add a whole new database backend.

In this exercise you will use Firely Server Facade libraries to build an ASP.NET Core library implementing a FHIR RESTful API on top of an existing database.

The existing database contains two simple tables 'Patient' and 'BloodPressure'. In the exercise we refer to it as the 'ViSi' system, short for 'VitalSigns'.

#### Git repository Vonk.Facade.Starter

This repository contains the completed exercise. You can find the repository at [Github](#). Of course we do recommend to try and do the exercise yourself first, before looking at the final result.

### 11.2.3 Prerequisites and Preparations

- Experience with programming ASP.NET (Core) libraries.
- Basic understanding of the [FHIR RESTful API](#) and FHIR servers.
- Visual Studio 2017 or newer
  1. get a free community edition at <https://www.visualstudio.com/downloads/>
  2. be sure to select the components for C# ASP.NET Core web development
- .NET Core 2.0 SDK, from <https://www.microsoft.com/net/download/windows>
  1. this is probably installed along with the latest Visual Studio, but needed if your VS is not up-to-date.
- SQL Server 2012 or newer:
  1. get a free developer or express edition at <https://www.microsoft.com/en-us/sql-server/sql-server-downloads>
  2. add SQL Server Management Studio from <https://docs.microsoft.com/en-us/sql/ssms/download-sql-server-management-studio-ssms>
- Postman, Fiddler or a similar tool to issue http requests and inspect the responses.

#### Installing the Firely Server

Before you can start implementing your facade, you need to have the Firely Server installed. See [Getting Started](#) on how to download the binaries and license key.

## Preparing the database

Download the `CreateDatabase` script, and create a SQL Server database with it.

It creates a database 'ViSi' with two tables: Patient and BloodPressure. You can familiarize yourself with the table structure and contents to prepare for the mapping to FHIR later on.

Proceed to the next step to start your facade project.

### 11.2.4 Starting your project

In this step, you will create and set up your project to start building the facade.

#### Create new project

1. Open Visual Studio 2017
2. File | New | Project
  - Choose Class Library (.NET Core)
  - Project name and directory at your liking; Click OK

#### Add Firely Server Packages

1. Tools > NuGet Package Manager > Package Manager Console
  - Run `Install-Package Vonk.Core`
  - Run `Install-Package Vonk.Fhir.R3` (if you want to use R3)
  - Run `Install-Package Vonk.Fhir.R4` (if you want to use R4)

---

**Note:** You can install the latest beta release of the Firely Server packages by adding `-IncludePrerelease` to the install command.

---

### 11.2.5 Mapping the database

In this step you will start mapping the existing database model to FHIR resources.

#### Reverse engineer the database model

To use [EF Core](#), install the package for the database provider(s) you want to target. This walkthrough uses SQL Server. For a list of available providers see Database Providers.

- Tools NuGet Package Manager Package Manager Console
- Run `Install-Package Microsoft.EntityFrameworkCore.SqlServer`

We will be using some Entity Framework Tools to create a model from the database. So we will install the tools package as well:

- Run `Install-Package Microsoft.EntityFrameworkCore.Tools`

**Note:** The current version of Firely Server uses the latest EF Core libraries. If you are developing for an older Firely Server version, please check the version of `Microsoft.EntityFrameworkCore.SqlServer.dll` in your Firely Server distribution folder. Add `-Version <version>` to the commands above to use the same version in your Facade implementation.

Now it's time to create the EF model based on your existing database.

- Tools NuGet Package Manager Package Manager Console
- Run the following command to create a model from the existing database. Adjust the Data Source to your instance of SQL Server. If you receive an error stating The term 'Scaffold-DbContext' is not recognized as the name of a cmdlet, then close and reopen Visual Studio.:

```
Scaffold-DbContext "MultipleActiveResultSets=true;Integrated Security=SSPI;Persist_
↪Security Info=False;Initial Catalog=ViSi;Data Source=localhost" Microsoft.
↪EntityFrameworkCore.SqlServer -OutputDir Models
//For localdb: Scaffold-DbContext "Server=(localdb)\mssqllocaldb;Database=ViSi;
↪Trusted_Connection=True;" Microsoft.EntityFrameworkCore.SqlServer -OutputDir_
↪Models
//For SQLEXPRESS: Scaffold-DbContext "Data Source=(local)\SQLEXPRESS;Initial_
↪Catalog=ViSi;Integrated Security=True" Microsoft.EntityFrameworkCore.SqlServer -
↪OutputDir Models
```

You can also generate the scaffolding using the [EF CLI tools](#) which are crossplatform:

```
dotnet ef dbcontext scaffold "User ID=SA;Password=<enter your password here>;
↪MultipleActiveResultSets=true;Server=tcp:.;Connect Timeout=5;Integrated Security=false;
↪Persist Security Info=False;Initial Catalog=ViSi;Data Source=localhost" Microsoft.
↪EntityFrameworkCore.SqlServer --output-dir Models
```

The reverse engineer process creates entity classes (`Patient.cs` & `BloodPressure.cs`) and a derived context (`ViSiContext.cs`) based on the schema of the existing database.

The entity classes are simple C# objects that represent the data you will be querying and saving. Later on you will use these classes to define your queries on and to map the resources from.

### Clean up generated code

- To avoid naming confusion with the FHIR Resourcetype Patient, rename both files and classes:
  - Patient ViSiPatient
  - BloodPressure ViSiBloodPressure

In `ViSiContext.cs`, ensure that the EF objects mapping our class to the database table are correct and without prefixes (since it's just our local classes that have them):

```
public virtual DbSet<ViSiBloodPressure> BloodPressure { get; set; }
public virtual DbSet<ViSiPatient> Patient { get; set; }
```

- The Scaffold command puts your connectionstring in the `ViSiContext` class. That is not very configurable. Later in the exercise, we will add it as 'DbOptions' to the `appsettings.instance.json` file in [3. Configure your Firely Server Facade](#).
  - Rename the default `Class1` class to `DbOptions`, and add this to interpret the setting:



```
public class DbOptions
{
    public string ConnectionString { get; set; }
}
```

- Remove the empty constructors from the ViSiContext class
- Use the options in your ViSiContext class, by adding:

```
private readonly IOptions<DbOptions> _dbOptionsAccessor;

public ViSiContext(IOptions<DbOptions> dbOptionsAccessor)
{
    _dbOptionsAccessor = dbOptionsAccessor;
}
```

- Change the existing OnConfiguring method that contains the connectionstring to:

```
protected override void OnConfiguring(DbContextOptionsBuilder optionsBuilder)
{
    if (!optionsBuilder.IsConfigured)
    {
        optionsBuilder.UseSqlServer(_dbOptionsAccessor.Value.ConnectionString);
    }
}
```

## Create your first mapping

1. Add a new public class ResourceMapper to the project
2. Add usings for Vonk.Core.Common, for Hl7.Fhir.Model, for Hl7.Fhir.Support and for <your project>.Models
3. Add a method to the class public IResource MapPatient(ViSiPatient source)
4. In this method, put code to create a FHIR Patient object, and fill its elements with data from the ViSiPatient:

```
var patient = new Patient
{
    Id = source.Id.ToString(),
    BirthDate = source.DateOfBirth.ToFhirDate()
};
patient.Identifier.Add(new Identifier("http://mycompany.org/patientnumber",
                                     source.PatientNumber));
// etc.
```

For more examples of filling the elements, see the FHIR API documentation: [FHIR-model](#).

5. Then return the created Patient object as an IResource with patient.ToIResource().

IResource is an abstraction from actual Resource objects as they are known to specific versions of the Hl7.Fhir.Net API. See [IResource](#).

## 11.2.6 Enable Search

Enabling search involves two major steps:

1. Creating a query to the database based on the bits and pieces in the search url
2. Getting a count and actual data from the database with that query, and map it to a `SearchResult`

The next paragraphs will walk you through these steps.

### 1. Create a query

Firely Server Facade is meant to be used across all kinds of database paradigms and schemas. Or even against underlying web services or stored procedures. This means Firely Server cannot prescribe the way your query should be expressed. After all, it could be an http call to a webservice, or a json command to MongoDB.

In our case we will build a LINQ query against our ViSi model, that is translated by Entity Framework to a SQL query. Because this is a quite common case, Firely Server provides a basis for it in the package `Vonk.Facade.Relational`.

- Go back to the NuGet Package Manager Console and run `Install-Package Vonk.Facade.Relational`

---

**Note:** If you did this previously for the other Firely Server packages, you can install the latest beta release of this package as well by adding `-IncludePrerelease` to the install command.

---

### Adding classes for the query

You usually create a query class per `ResourceType`. The `Query` object is used to capture the elements of the search that are provided to the `QueryFactory`.

In this exercise we start with resource type `Patient`, and will create a `PatientQuery` and `PatientQueryFactory` class. Because `PatientQuery` has no specific content of its own, we will include both in one file.

- Add a new class `PatientQueryFactory` to the root of the project
- Add using statements for `Vonk.Facade.Relational`, `Microsoft.EntityFrameworkCore`, and `<your project>.Models`
- Above the actual `PatientQueryFactory` class insert the `PatientQuery` class:

```
public class PatientQuery: RelationalQuery<ViSiPatient>
{}
```

- Now flesh out the `PatientQueryFactory`:

```
public class PatientQueryFactory: RelationalQueryFactory<ViSiPatient, PatientQuery>
{}
```

## Adding a constructor

You have to provide a constructor for the factory class. With this you tell Firely Server for which resource type this QueryFactory is valid. The DbContext is used for retrieving DbSetes for related entities, as we will see later:

```
public PatientQueryFactory(DbContext onContext) : base("Patient", onContext) { }
```

## Deciding on a FHIR version

You need to explicitly tell Firely Server for which FHIR version(s) you wish to return resources. If you don't override EntryInformationModel, any search will fail with a 501 Not Implemented. The following override will allow searches for any possible FHIR version to be handled by your facade:

```
public override PatientQuery EntryInformationModel(string informationModel)
{
    return default(PatientQuery);
}
```

If you wish to implement search only for a single FHIR version or for a limited set of versions you can override the method like this:

```
public override PatientQuery EntryInformationModel(string informationModel)
{
    if (informationModel == VonkConstants.Model.FhirR4)
    {
        return default(PatientQuery);
    }

    throw new NotImplementedException($"FHIR version {informationModel} is not supported
    ↪");
}
```

## Handling the search request

Each of the searchparameters in the search request triggers a call to the Filter method. This method takes a parameterName and IFilterValue as its arguments.

The parameterName is the searchparameter as it was used in the search url. This name corresponds with the code field in a SearchParameter resource. The IFilterValue value is one of 10 possible implementations, one for each type of SearchParameter. See *IFilterValue implementations* for a short description of these possibilities.

By default the Filter method dispatches the call to a suitable overload of AddValueFilter, based on the actual type of the value parameter. It is up to you to override the ones you support any parameters for.

- Override the method PatientQuery AddValueFilter(string parameterName, TokenValue value) in the PatientQueryFactory class to implement support for the \_id parameter, which is a token type parameter.

The \_id parameter must be matched against the ViSiPatient.Id property. So we have to:

- Parse the Token.Code to an integer (ViSiPatient.Id is of type int)
- Create a query with a predicate on ViSiPatient.Id.

This is how:

```

if (parameterName == "_id")
{
    if (!int.TryParse(value.Code, out int patientId))
    {
        throw new ArgumentException("Patient Id must be an integer value.");
    }
    else
    {
        return PredicateQuery(vp => vp.Id == patientId);
    }
}
return base.AddValueFilter(parameterName, value);

```

**Note:** The `ArgumentException` in this code will automatically result in setting the argument status to error, so the Firely Server will send a response with an error code and `OperationOutcome`. See the information about the `IArgumentCollection` and `IArgument` classes in *IVonkContext*.

That's it for now, we will add support for another parameter later.

## IFilterValue implementations

There are 10 possible implementations you can use as value for the `IFilterValue` parameter in the Query. The first 7 are the *general search parameter types*: `StringValue`, `DateTimeValue`, `TokenValue`, `NumberValue`, `QuantityValue`, `UriValue` and `ReferenceValue`.

Besides that there are two special values for chaining and reverse chaining: `ReferenceToValue` and `ReferenceFromValue`.

And finally there is a special value for when Firely Server does not know the `SearchParameter` and hence not the type of it: `RawValue`.

This is the second step for enabling search.

### 11.2.7 2. Get the data and map to FHIR

Getting the data happens in the implementation of the `ISearchRepository`. It has only one method, `Search`. The `Vonk.Facade.Relational` package has an abstract implementation of it that you can use as a starting point. This implementation assumes that you can support searching for exactly one `ResourceType` at a time. The gist of the implementation is to switch the querying based on the `ResourceType`. The querying itself then looks pretty much the same for every type of resource.

- Add the new class `ViSiRepository` to the root of the project:

```
public class ViSiRepository : SearchRepository
```

- You have to provide a constructor that gets a `QueryContext`. We'll get to that later. Apart from that you will need your `DbContext` (`ViSiContext`) to query on, and the `ResourceMapper` to perform the mapping of the results. So put all of that in the constructor:

```
private readonly ViSiContext _visiContext;
private readonly ResourceMapper _resourceMapper;
```

(continues on next page)

(continued from previous page)

```

public ViSiRepository(QueryContext queryContext, ViSiContext visiContext,
↳ ResourceMapper resourceMapper) : base(queryContext)
{
    _visiContext = visiContext;
    _resourceMapper = resourceMapper;
}

```

- You will have to implement the abstract method `Task<SearchResult> Search(string resourceType, IArgumentCollection arguments, SearchOptions options)`.

- First, let's inspect the parameters:

#### resourceType

The ResourceType that is being searched for, e.g. Patient in `<firely-server-endpoint>/Patient?...`

#### arguments

All the arguments provided in the search, whether they come from the path (like 'Patient'), the querystring (after the '?'), the headers or the body. Usually you don't have to inspect these yourself.

#### options

A few hints on how the query should be executed: are deleted or contained resources allowed etc. Usually you just pass these on as well.

- The pattern of the implementation is:

1. switch on the resourceType
2. dispatch to a method for querying for that resourceType

Naturally we do this async, since in a web application you should never block a thread while waiting for the database.

- To implement this, add this to the class:

```

protected override async Task<SearchResult> Search(string resourceType,
↳ IArgumentCollection arguments, SearchOptions options)
{
    switch (resourceType)
    {
        case "Patient":
            return await SearchPatient(arguments, options);
        default:
            throw new NotImplementedException($"ResourceType {resourceType} is
↳ not supported.");
    }
}

```

- Now we moved the problem to `SearchPatient`, so this method needs to be implemented. The pattern here is:

1. Create a query - in this case, a `PatientQuery` via `PatientQueryFactory`.
2. Execute the query against the `DbContext` (our `_visiContext`) to get a count of matches.
3. Execute the query against the `DbContext` to get the current page of results.
4. Map the results using the `_resourceMapper`

The implementation of this looks like:

```
private async Task<SearchResult> SearchPatient(IArgumentCollection arguments, SearchOptions options)
{
    var query = _queryContext.CreateQuery(new PatientQueryFactory(_visiContext), arguments, options);

    var count = await query.ExecuteCount(_visiContext);
    var patientResources = new List<IResource>();

    if (count > 0)
    {
        var visiPatients = await query.Execute(_visiContext).ToListAsync();

        foreach (var visiPatient in visiPatients)
        {
            patientResources.Add(_resourceMapper.MapPatient(visiPatient));
        }
    }
    return new SearchResult(patientResources, query.GetPageSize(), count);
}
```

What happens behind the scenes is that the `QueryBuilderContext` creates a `QueryBuilder` that analyzes all the arguments and options, and translates that into calls into your `PatientQueryFactory`. This pattern offers maximum assistance in processing the search, but also gives you full control over the raw arguments in case you need that for anything.

---

### Throwing exceptions during query building

When calling `QueryContext.CreateQuery(...)`, the underlying code (`QueryBuilder` and `ArgumentHandler`) uses your `PatientQueryFactory` to create a `PatientQuery` to retrieve the data.

Both `QueryBuilder` and `ArgumentHandler` are part of `Vonk.Core` and follow the default behavior of Firely Server.

One consequence is that, if you throw exceptions in `PatientQueryFactory`, they will be converted by `ArgumentHandler` to an `Issue` on the argument, and the argument will be ignored. These arguments are later automatically converted to an `OperationOutcome` in the search result bundle to tell the caller he made a mistake.

If you want to act on arguments with problems, you can do so by inspecting the argument collection after query creation, in your implementation of `ISearchRepository`.

---

In the next paragraph you will configure your Firely Server to use your Facade, and can – finally – try out some searches. The paragraph after that expands the project to support `ViSiBloodPressure` Observations, and details how to add custom search parameters.

---

**Note:** Your implementation of `ISearchRepository` may have other dependencies than the ones listed above, but it cannot be dependent upon `IStructureDefinitionSummaryProvider`. That causes a circular dependency, and will have you wait for a response from the server indefinitely. This means that in the implementation you can work with POCO's (as is done in this tutorial) or with 'raw' `SourceNode` instances, but not with `ITypedElement` (for the latter two see the Firely .NET SDK).

---

## 11.2.8 Finalizing search

In the previous steps you have created search support for the `_id` parameter on a Patient resource type. In order to test if your Facade implementation works correctly, you will need to perform a couple of steps:

1. Create a configuration class for the [ASP .Net Core pipeline](#)
2. Plug the Facade into the Firely Server
3. Configure the Firely Server to use your repository

### 1. Add configuration class

To add your repository service to the Firely Server pipeline, you will need to add a configuration class that sets the order of inclusion, and adds to the services. For background information, see [Configuration classes](#).

- Add a static class to your project called `ViSiConfiguration`
- Add the following code to it:

```
[VonkConfiguration(order: 210)]
public static class ViSiConfiguration
{
    public static IServiceCollection AddViSiServices(this IServiceCollection
↪services, IConfiguration configuration)
    {
        services.AddDbContext<ViSiContext>();
        services.TryAddSingleton<ResourceMapper>();
        services.TryAddScoped<ISearchRepository, ViSiRepository>();

        services.Configure<DbOptions>(configuration.GetSection(nameof(DbOptions)));
        return services;
    }
}
```

**Note:** As of Firely Server version 4.4.0, your Facade should not have an order greater than or equal to 211. The reason for this is that upon configuring the administration database, Firely Server checks whether an `ISearchRepository` is registered. The earliest of these configurations is at order 211.

**Note:** Should you need to register any other classes or interface, please see [Register a service in your plugin](#) for more background.

### 2. Create your Facade plugin

- First, build your project
- Find the resulting dll and copy that to the `plugins` folder in the working directory of your Firely Server

**Note:** If your Firely Server working directory does not contain a `plugins` folder yet, you can create one. Within it, you can create subfolders, which can be useful if you work with multiple plugins.

You can also configure the name and location of this folder with the `PipelineOptions.PluginDirectory` setting in the `appsettings` file.

---

### 3. Configure your Firely Server Facade

- Create an `appsettings.instance.json` file in your Firely Server working directory.

**Tip:** See *Firely Server settings* for more information about the hierarchy of the `appsettings(.*)`.json files and the settings that can be configured.

---

- Add a setting for the `connectionstring` to the `appsettings.instance.json` file:

```
"DbOptions" : { "ConnectionString" : "<paste the connection string to your ViSi_
↪database here>" },
```

- Add the `SupportedInteractions` section. You can look at *Enable or disable interactions* to check what this section should contain. For now you only need `"WholeSystemInteractions": "capabilities"`, `"InstanceLevelInteractions": "read"` and `"TypeLevelInteractions": "search"`:

```
"SupportedInteractions": {
  "InstanceLevelInteractions": "read",
  "TypeLevelInteractions": "search",
  "WholeSystemInteractions": "capabilities"
},
```

- Add the `SupportedModel` section to indicate which resource types and search parameters you support in your Facade implementation:

```
"SupportedModel": {
  "RestrictToResources": [ "Patient" ],
  "RestrictToSearchParameters": ["Resource._id", "StructureDefinition.url"]
},
```

- You will need to add your repository to the Firely Server pipeline, and remove the existing repository implementations. The standard settings for the pipeline configuration can be found in the `appsettings.default.json` file, or see *Configure the pipeline* for an example.

- Copy the whole `PipelineOptions` section to your `appsettings.instance.json` file (both `/` and `/administration`)
- To the `Include` part of the branch with `"Path":"/"` add your namespace, and remove the `Vonk.Repository.*` lines from it:

```
{
  "Path": "/",
  "Include": [
    "Vonk.Core",
    "Vonk.Fhir.R3",
    "Vonk.Subscriptions",
    "Vonk.Smart",
    "Vonk.UI.Demo",
    "ViSiProject" // fill in (a prefix of) the namespace of your project here
  ]
}
```

(continues on next page)



(continued from previous page)

```
]
},
```

- Remove the PipelineOptions from appsettings.default.json, because of the warning mentioned on the [Hierarchy of settings](#).

## Test your work

Proceed to the next section to test your Facade, and for some helpful tips about debugging your code.

### 11.2.9 Debugging the Facade

- Start your Firely Server

---

**Note:** If this is your first startup of Firely Server, it will take a while to load in all of the specification files.

---

- You can inspect the console log to see if the pipeline is configured to include your repository. See [Detailed logging of loading plugins](#) for more details.
- To test your Facade, open Postman, or Fiddler, or use curl to request GET `http://localhost:4080/metadata`. The resulting CapabilityStatement should list only the Patient resource type in its .rest.resource field, and – among others – the \_id search parameter in the .rest.searchParam field.
- Now you can test that searching patients by \_id works: GET `http://localhost:4080/Patient?_id=1`. Requesting the resource ‘normally’ should automatically work as well: GET `http://localhost:4080/Patient/1`

---

**Important:** If it works, congratulations! You now have a Firely Server Facade running!

---

## Testing during implementation

Follow these steps if you want to test your work during the implementation phase without having to build, copy and start Firely Server each time, or with the ability to set break points in your code and debugging it:

- In the project properties, click on the Build tab.
- Set the Output path to your Firely Server plugins directory.
- Go to the Debug tab and set Launch to Executable.
- Point the Executable field to your dotnet.exe.
- Set the Application arguments to `<your-Firely-Server-working-directory>/Firely.Server.dll`.
- Set the Working directory to your Firely Server working directory.

Now, whenever you click to start debugging, Firely Server will start from your project and your project dll will be automatically built to the Firely Server plugins directory.

## Next part of the exercise

You can proceed to the next section to add support for Observations as well.

### 11.2.10 Finalizing your project

#### Add support for the ViSiBloodPressure Observations

First, follow similar steps as above to support ViSiBloodPressure:

1. Add a mapping method in the `ResourceMapper` class to map from a `ViSiBloodPressure` to a FHIR Observation resource, and return that as an `IResource`.
2. Create a `BloodPressureQuery` query class.
3. Add a `BPQueryFactory` extending `RelationalQueryFactory<ViSiBloodPressure, BloodPressureQuery>`.
4. Implement support for the `_id` parameter by overriding `public virtual BloodPressureQuery AddValueFilter(string parameterName, TokenValue value)`.
5. Add the Observation type to the `SupportedModel` section in Firely Server's `appsettings.instance.json`:  
"RestrictToResources": [ "Patient", "Observation" ]

When you have completed these steps, build your project again and copy the dll to your Firely Server plugins folder. After you (re)start Firely Server, you will be able to request an Observation through your Facade: GET `http://localhost:4080/Observation?_id=1` or GET `http://localhost:4080/Observation/1`.

Since you do not always want to request Observations by their technical id, but more often might want to request Observations from a specific patient, the next part will describe implementing support that. The Patient resource is referenced by the Observation in its subject field. The corresponding search parameter is either `subject` or `patient`.

#### Add support for chaining

To add support for searching on `Observation?subject:Patient._id` we need to override the `AddValueFilter` overload receiving a `ReferenceToValue` parameter in the query factory for BloodPressure (`BPQueryFactory`).

The `ReferenceToValue` type contains the possible `Targets` for the chain search parameter as parsed from the query string. We are currently interested only on the `Patient` type so we can restrict the implementation to that target. The `ReferenceToValue` type also has an extension method `CreateQuery` that expects an implementation of the `RelationalQueryFactory` of the referenced target. This will generate the query to obtain the resources referenced by it.

Searching on chained parameters involves the following steps:

1. Retrieve all patient ids based on the chained parameter. You can use the `ReferenceToValue.CreateQuery` extension method to get the query and run the query with its `Execute` method.
2. Create a `PredicateQuery` with the condition that `ViSiBloodPressure.PatientId` is included in the ids retrieved at the previous step.

The final code should look similar to this:

```
public override BloodPressureQuery AddValueFilter(string parameterName,
    ↪ReferenceToValue value)
{
    if (parameterName == "subject" && value.Targets.Contains("Patient"))
```

(continues on next page)

(continued from previous page)

```

    {
        var patientQuery = value.CreateQuery(new
        PatientQueryFactory(OnContext));
        var patIds = patientQuery.Execute(OnContext).Select(p => p.Id);

        return PredicateQuery(bp => patIds.Contains(bp.PatientId));
    }
    return base.AddValueFilter(parameterName, value);
}

```

**Note:** patIds is of type IQueryable, so the resulting BloodPressureQuery will still be executed as a single command to the database.

3. Add support for the Observation.subject search parameter in the Firely Server appsettings similar to how we did it for \_id.

At this point you should be able to search for GET [http://localhost:4080/Observation?subject:Patient.\\_id=1](http://localhost:4080/Observation?subject:Patient._id=1)

### Add support for reverse chaining

Adding support for Patient?\_has:Observation:subject:\_id=1 is similar. You just need to use the AddValueFilter overload receiving a ReferenceFromValue.

The ReferenceFromValue type has a Source property filled in with the source of the search parameter. It also has an extension method CreateQuery that given the corresponding RelationalQueryFactory implementation can generate the query to obtain resources referenced by the reverse chaining.

So you can add reverse chaining with the following code:

```

public class PatientQueryFactory
{
    public override PatientQuery AddValueFilter(string parameterName, ReferenceFromValue
    value)
    {
        if (parameterName == "subject" && value.Source == "Observation")
        {
            var obsQuery = value.CreateQuery(new BPQueryFactory(OnContext));
            var obsIds = obsQuery.Execute(OnContext).Select(bp => bp.PatientId);

            return PredicateQuery(p => obsIds.Contains(p.Id));
        }
        return base.AddValueFilter(parameterName, value);
    }
}

```

**Note:** The reverse chaining example above uses the Patient resource type as its base, so you will need to implement this in your PatientQueryFactory.

Now you can test if reverse chaining works: `http://localhost:4080/Patient?_has:Observation:subject:_id=1`

### Get the goodies

At this point you get out of the box support for `_include` and `_revinclude` (:iterate as well), and combinations of search parameters. You can test the following scenarios:

1. `_include`: `http://localhost:4080/Observation?_include=Observation:subject`
2. `_revinclude`: `http://localhost:4080/Patient?_revinclude=Observation:subject`
3. combinations of the above

### Adding a custom SearchParameter

Your Firely Server will load the standard parameters from the specification on first startup, so the `_id` SearchParameter from the exercise is already known to Firely Server, as well as any of the other standard search parameters for the resource types.

If you want to implement support for a custom search parameter, you will need to have the definition of that in the form of a SearchParameter resource, and add it to your Firely Server. The [Configure Search Parameters](#) section describes how to do that.

Of course you will also need to implement the correct `AddValueFilter` method in your `<resourcetype>QueryFactory` to handle the parameter correctly, as is done for the `_id` parameter in the exercise.

### The end?

This concludes the exercise. An example [Github repository](#) contains the completed exercise.

Please feel free to try out more options, and [ask for help](#) if you get stuck!

The next topic will show you how to enable *Create, Update and Delete* interactions.

### Postscript

If your resource is split across multiple tables in the database, you'll need to make use of `.Include()` to have EF load the dependent table. To do so in Firely Server, override the `GetEntitySet()` method in your *RelationalQuery* class, for example

```
protected override IQueryable<ViSiPatient> GetEntitySet(DbContext dbContext)
{
    // load the dependent Address table
    return dbContext.Set<ViSiPatient>().Include(p => p.Address).AsNoTracking();
}
```

### 11.2.11 Enable changes to the repository

In *Exercise: Build your first Facade* you have created read and search support for a Firely Server Facade on top of an existing database. The next part will walk you through enabling create, update and delete support. This will be done in three steps:

1. Map the FHIR resource to the database model;
2. Implement the `IResourceChangeRepository`;
3. Indicate support in the appsettings file

#### 1. Map the FHIR data to the model

For both the Create and Update interactions, we need to map our incoming FHIR data to the ViSi model. To do that, we will add new mapping methods to our `ResourceMapper` class.

- Add a method called `MapViSiPatient` to the `ResourceMapper` class, that takes a FHIR `Patient` and returns a `ViSiPatient`.
- Implement the mapping from the FHIR object to the `ViSiPatient` model:

```
public ViSiPatient MapViSiPatient(IResource source)
{
    var fhirPatient = source.ToPoco<Patient>();
    var visiPatient = new ViSiPatient();

    if (source.Id != null)
    {
        if (!int.TryParse(source.Id, out int id))
            throw new VonkRepositoryException("Id needs to be integer to map_
↪resource");
        visiPatient.Id = id;
    }
    visiPatient.PatientNumber = fhirPatient.Identifier.Find(i =>
        (i.System == "http://mycompany.org/patientnumber")).
↪Value;

    // etc.

    return visiPatient;
}
```

- Where it says 'etc.', fill in the rest of the code to map the data to required fields of the database, and any other fields you have data for.

## 2. Implement the IResourceChangeRepository

You are going to implement a repository that handles changes to your database. The interface for this is called `IResourceChangeRepository`, which can be found in `Vonk.Core.Repository`.

- Add a new class `ViSiChangeRepository` to the project, that implements the `IResourceChangeRepository`:

```
public class ViSiChangeRepository : IResourceChangeRepository
```

- Choose to implement the interface, so the required methods are added to the class.
- Just like with the search repository, you will need your `DbContext` to query on, and the `ResourceMapper` to perform the mapping of the incoming data to your proprietary model.

So put all of that in the constructor:

```
private readonly ViSiContext _visiContext;
private readonly ResourceMapper _resourceMapper;

public ViSiChangeRepository(ViSiContext visiContext, ResourceMapper resourceMapper)
{
    _visiContext = visiContext;
    _resourceMapper = resourceMapper;
}
```

### Implementing Create

- Now implement the `Create` method with a switch on resource type, so you can add other resource types later:

```
public async Task<IResource> Create(IResource input)
{
    switch (input.Type)
    {
        case "Patient":
            return await CreatePatient(input);
        default:
            throw new NotImplementedException($"ResourceType {input.Type} is not_
↪supported.");
    }
}
```

- As you can see, we have deferred the work to a `CreatePatient` method, which we also need to implement. This method will add the new resource to the collection, and save the changes to the database:

```
private async Task<IResource> CreatePatient(IResource input)
{
    var visiPatient = _resourceMapper.MapViSiPatient(input);

    await _visiContext.Patient.AddAsync(visiPatient);
    await _visiContext.SaveChangesAsync();

    // return the new resource as it was stored by this server
    return _resourceMapper.MapPatient(_visiContext.Patient.Last());
}
```

- For the Create and Update methods, you will also need to implement the `NewId` and `NewVersion` methods, because Firely Server will call them. For the `NewId` method, we will return null, since our ViSi database does not allow us to create our own index value. Since our ViSi repository does not handle versions, we will let the `NewVersion` method return null as well:

```
public string NewId(string resourceType)
{
    return null;
}

public string NewVersion(string resourceType, string resourceId)
{
    return null;
}
```

**Note:** For the ViSi repository we're using a null value, but you can implement this method any way that's useful for your own repository. The public Firely Server for example generates a GUID in these methods.

At this point you can skip ahead to [3. Configure the service and Firely Server](#), if you want to try and create a new patient in the ViSi database.

**Tip:** This is easiest to test if you retrieve an existing resource from the database first with your HTTP tool. Then change some of the data in the resulting JSON or XML, and send that back to your Facade.

## Implementing Update

Implementing the Update method can be done like the Create, with a switch on resource type, and instead of adding a resource to the collection, you will update the collection:

```
private async Task<IResource> UpdatePatient(ResourceKey original, IResource update)
{
    try
    {
        var visiPatient = _resourceMapper.MapViSiPatient(update);

        var result = _visiContext.Patient.Update(visiPatient);
        await _visiContext.SaveChangesAsync();

        return _resourceMapper.MapPatient(result.Entity);
    }
    catch (Exception ex)
    {
        throw new VonkRepositoryException($"Error on update of {original} to {update.
        ↪Key()}", ex);
    }
}
```

## Implementing Delete

Deleting a resource from the collection is done by first looking up the corresponding resource, and then removing it from the collection. Note that the database used for this exercise cannot process the deletion of the Patient when there are still related Observations in the BloodPressure table, so we need to remove them as well or choose to throw an error.

- First, create a switch on resource type in the main Delete method again.
- Implement the DeletePatient:

```
private async Task<IResource> DeletePatient(ResourceKey toDelete)
{
    int toDelete_id = int.Parse(toDelete.ResourceId);
    var visiPatient = _visiContext.Patient.Find(toDelete_id);

    var bpEntries = _visiContext.BloodPressure.Where(bp => bp.PatientId == toDelete_
↪id);

    var result = _resourceMapper.MapPatient(visiPatient);

    try
    {
        _visiContext.BloodPressure.RemoveRange(bpEntries);
        _visiContext.Patient.Remove(visiPatient);
        await _visiContext.SaveChangesAsync();
    }
    catch (Exception ex)
    {
        throw new VonkRepositoryException($"Error on deleting Patient with Id
↪{toDelete_id}", ex);
    }

    return result;
}
```

## 3. Configure the service and Firely Server

Just like with the search repository, you will need to add your change repository as service to the pipeline. Also, you will need to indicate support for the CRUD interactions in your Firely Server appsettings.

- In your project, go to the ViSiConfiguration class, and add this line to add an IResourceChangeRepository to the pipeline:

```
services.TryAddScoped<IResourceChangeRepository, ViSiChangeRepository>();
```

- Add support for the interactions to the SupportedModel section of the Firely Server appsettings:

```
"SupportedInteractions": {
  "InstanceLevelInteractions": "read, update, delete",
  "TypeLevelInteractions": "search, create",
  "WholeSystemInteractions": "capabilities, search"
},
```

- Adjust PipelineOptions.Branches.Include from Vonk.Core.Operations.Crud.ReadConfiguration to ``Vonk.Core.Operations include all operations, including Create.



You can now build your project, copy the dll to the Firely Server plugins folder and run Firely Server to test the new interactions on your Facade.

## The end?

This concludes the second exercise. Please feel free to try out more options, and *ask for help* if you get stuck!

The next topic will show you how to integrate *Access Control*.

## 11.3 Plugins available for Firely Server

### 11.3.1 Infrastructural plugins

#### Name

Scheduler

#### Configuration

Vonk.Core.Quartz.QuartzConfiguration

#### License token

<http://fire.ly/vonk/plugins/infra>

#### Order

10

#### Description

Registers a scheduler that can run jobs periodically. You can use this yourself, but with care (you don't want jobs slowing down the server):

- Implement a Quartz.IJob, let's say with class MyJob { ... }
- Have the Quartz.IScheduler injected
- Call `IScheduler.StartJob<MyJob>(TimeSpan runInterval, CancellationToken cancellationToken)`

#### Name

Maintenance

#### Configuration

Vonk.Core.Infra.MaintenanceConfiguration

#### License token

<http://fire.ly/vonk/plugins/infra>

#### Order

20

#### Description

Periodically cleans the indexed values for deleted or superseded resources from the database.

#### Name

License

#### Configuration

Vonk.Core.Licensing.LicenseConfiguration

#### License token

<http://fire.ly/vonk/plugins/infra>

**Order**

120

**Description**

Registers the LicenseService that checks for a valid license. Without this plugin Firely Server does not work.

**Name**

Serialization

**Configuration**

Vonk.Core.Serialization.SerializationConfiguration

**License token**<http://fire.ly/vonk/plugins/infra>**Order**

130

**Description**

Registers an implementation for the ISerializationService and ISerializationSupport interfaces and actual serializers and parsers for JSON and XML.

**Name**

Pluggability

**Configuration**

Vonk.Core.Pluggability.PluggabilityConfiguration

**License token**<http://fire.ly/vonk/plugins/infra>**Order**

150

**Description**

Registers services to dynamically build the IModelService using registered IModelContributor implementations.

**Name**

Http to Vonk

**Configuration**

Vonk.Core.Context.Http.HttpToVonkConfiguration

**License token**<http://fire.ly/vonk/plugins/http>**Order**

1110

**Description**

Builds an *IVonkContext* out of the *HttpContext*. You can only access the IVonkContext in the pipeline from plugins that have a higher order.

**Name**

Vonk to Http

**Configuration**

Vonk.Core.Context.Http.VonkToHttpConfiguration

**License token**<http://fire.ly/vonk/plugins/http>

**Order**

1120

**Description**

Translates the response in the *IVonkContext* to a response on the *HttpContext*. It honors the value of the *prefer* header if present. It also adds the *VonkExceptionHandler* to the pipeline as a last resort for catching exceptions.

**Name**

Formatter

**Configuration**`Vonk.Core.Context.Format.FormatConfiguration`**License token**<http://fire.ly/vonk/plugins/infra>**Order**

1130

**Description**

Registers an implementation of *IFormatter* that can write the *IVonkContext.Response.Payload* to the response body in the requested format. Does not add a processor to the pipeline.

**Name**

Liveness probe

**Configuration**`Vonk.Core.Infra.LivenessCheckConfiguration`**License token**

&lt;none&gt;

**Order**

1160

**Description**

Check whether the server is active, for use in e.g. Kubernetes liveness probes or Docker healthchecks. See *Liveness and readiness* for possible results.

**Name**

Long running tasks

**Configuration**`Vonk.Core.Infra.LongRunning.LongRunningConfiguration`**License token**<http://fire.ly/vonk/plugins/infra>**Order**

1170

**Description**

If Vonk processes a task that could lead to inconsistent output, all other requests are rejected by this plugin. Long running tasks are e.g. the *Import of Conformance Resources* and *Re-indexing for new or changed SearchParameters*.

**Name**

Readiness probe

**Configuration**`Vonk.Core.Infra.ReadinessCheckConfiguration`

**License token**

<none>

**Order**

1180

**Description**

Check whether the server can handle requests, for use in e.g. Kubernetes readiness probes or Docker healthchecks. See *Liveness and readiness* for possible results.

**Name**

Compartments

**Configuration**

`Vonk.Core.Context.Features.CompartmentsConfiguration`

**License token**

<http://fire.ly/vonk/plugins/search>

**Order**

1210

**Description**

Recognizes a compartment in a compartment search on system or type level (see *Search*). It is added as a feature of type `ICompartment` to the `IVonkContext.Features` collection, to be used by *Search* later on. This `ICompartment` feature will limit all queries to within the specified compartment.

**Name**

Supported Interactions

**Configuration**

`Vonk.Core.Context.Guards.SupportedInteractionsConfiguration`

**License token**

<http://fire.ly/vonk/plugins/infra>

**Order**

1220

**Description**

Blocks interactions that are not listed as supported.

**Options**

`SupportedInteractions`, see *Enable or disable interactions*.

**Name**

Size Limits

**Configuration**

`Vonk.Core.Context.Guards.SizeLimitsConfiguration`

**License token**

<http://fire.ly/vonk/plugins/infra>

**Order**

1225

**Description**

Rejects bodies that are too large and bundles with too many entries.

**Options**

`SizeLimits`, see *Protect against large input*

**Name**

Url mapping

**Configuration**`Vonk.Core.Context.UrlMapping.UrlMappingConfiguration`**License token**<http://fire.ly/vonk/plugins/infra>**Order**

1235

**Description**

In a resource in the request, urls pointing to this instance of Firely Server are made relative. In a resource in the response, relative urls are made absolute, by adding the base url of the server. This way the server can be addressed in multiple ways (e.g. <http://intranet.acme.com/fhir> and <https://fhir.acme.com>) and still provide correct absolute urls.

**Name**

Default Shapes

**Configuration**`Vonk.Core.Context.Guards.DefaultShapesConfiguration`**License token**<http://fire.ly/vonk/plugins/infra>**Order**

4110

**Description**

If no sort order is given for a search, `_lastUpdated:asc` is added. If no count is given for a search, `_count=<default count>` is added.

**Options**

`BundleOptions.DefaultCount`, see [Search size](#).

### 11.3.2 Support for different FHIR versions

**Name**

FHIR R3

**Configuration**`Vonk.Fhir.R3.FhirR3Configuration`**License token**<http://fire.ly/vonk/plugins/fhirr3>**Order**

100

**Description**

Registers services to support FHIR STU3 (or R3).

**Name**

FHIR R3 Specification

**Configuration**`Vonk.Fhir.R3.FhirR3SpecificationConfiguration`

**License token**

<http://fire.ly/vonk/plugins/fhrr3>

**Order**

112

**Description**

Registers an `Hl7.Fhir.Specification.IStructureDefinitionSummaryProvider` for FHIR STU3 (or R3).

**Name**

FHIR R3 Validation

**Configuration**

`Vonk.Fhir.R3.Validation.ValidationConfigurationR3`

**License token**

<http://fire.ly/vonk/plugins/fhrr3>

**Order**

4845

**Description**

Registers a validator for FHIR STU3 (or R3).

**Name**

FHIR R4

**Configuration**

`Vonk.Fhir.R4.FhirR4Configuration`

**License token**

<http://fire.ly/vonk/plugins/fhrr4>

**Order**

101

**Description**

Registers services to support FHIR R4.

**Name**

FHIR R4 Specification

**Configuration**

`Vonk.Fhir.R4.FhirR4SpecificationConfiguration`

**License token**

<http://fire.ly/vonk/plugins/fhrr4>

**Order**

112

**Description**

Registers an `Hl7.Fhir.Specification.IStructureDefinitionSummaryProvider` for FHIR R4.

**Name**

FHIR R4 Validation

**Configuration**

`Vonk.Fhir.R4.Validation.ValidationConfigurationR4`

**License token**

<http://fire.ly/vonk/plugins/fhrr4>

**Order**

4845

**Description**

Registers a validator for FHIR R4.

### 11.3.3 FHIR RESTful interactions

**Name**

Read

**Configuration**`Vonk.Core.Operations.Crud.ReadConfiguration`**License token**<http://fire.ly/vonk/plugins/read>**Order**

4410

**Description**

Implements FHIR instance read. It will return the Resource that matches the id *and* the FHIR version.  
 If a Resource with matching id is found with another FHIR version you are notified.

**Name**

Create

**Configuration**`Vonk.Core.Operations.Crud.CreateConfiguration`**License token**<http://fire.ly/vonk/plugins/create>**Order**

4420

**Description**

Implements FHIR type create.

**Name**

Update

**Configuration**`Vonk.Core.Operations.Crud.UpdateConfiguration`**License token**<http://fire.ly/vonk/plugins/update>**Order**

4430

**Description**

Implements FHIR instance update, with support for 'upsert': creating a Resource with a pre-assigned id. Note that id's must be unique across FHIR versions.

**Name**

Patch

**Configuration**`Vonk.Core.Operations.Crud.FhirPatchConfiguration`

**License token**

<http://fire.ly/vonk/plugins/update>

**Order**

4433

**Description**

Implements FHIR instance patch, as specified by [FHIR Patch](#).

**Name**

Delete

**Configuration**

`Vonk.Core.Operations.Crud.DeleteConfiguration`

**License token**

<http://fire.ly/vonk/plugins/delete>

**Order**

4440

**Description**

Implements FHIR instance delete. Since id's in Firely Server must be unique across FHIR versions, the delete is issued on the provided id, regardless of the FHIR version.

**Name**

Search

**Configuration**

`Vonk.Core.Operations.Search.SearchConfiguration`

**License token**

<http://fire.ly/vonk/plugins/search>

**Description**

Implements FHIR Search on system and type level. For data access it uses the registered implementation of `ISearchRepository`, which can be any of the implementations provided by Firely Server or an implementation provided by a Facade plugin. The implementations provided by Firely Server also require the Index plugin to extract searchparameter values from the resources.

**Order**

4220

**Options**

- `AdministrationImportOptions`, see *SearchParameters and other Conformance Resources*, for available Searchparameters
- `SupportedModel.RestrictToSearchParameters`, see *Restrict supported resources and SearchParameters* for available Searchparameters
- `BundleOptions`, see *Search size*, for number of returned results

See *ISearchRepository* and *Firely Server Facade*.

**Name**

Search support

**Configuration**

`Vonk.Core.Repository.RepositorySearchSupportConfiguration`

**License token**

<http://fire.ly/vonk/plugins/search>



**Order**

140

**Description**

Registers services required for Search. It is automatically registered by Search.

**Name**

Index

**Configuration**

```
Vonk.Core.Repository.RepositoryIndexSupportConfiguration
```

**License token**

<http://fire.ly/vonk/plugins/index>

**Order**

141

**Description**

Extracts values matching Searchparameters from resources, so they can be searched on.

**Name**

Include

**Configuration**

```
Vonk.Core.Operations.Search.IncludeConfiguration
```

**License token**

<http://fire.ly/vonk/plugins/include>

**Order**

4210

**Description**

Implements `_include` and `_revinclude`. This acts on the result bundle of a search. Therefore it also works out of the box for Facade implementations, provided that the Facade implements support for the reference Searchparameters that are used in the `_(rev)include`.

**Name**

Elements

**Configuration**

```
Vonk.Core.Context.Elements.ElementsConfiguration
```

**License token**

<http://fire.ly/vonk/plugins/search>

**Order**

1240

**Description**

Applies the `_elements` parameter to the Resource that is in the response (single resource or bundle).

**Name**

Summary

**Configuration**

```
Vonk.Core.Context.Elements.SummaryConfiguration
```

**License token**

<http://fire.ly/vonk/plugins/search>

**Order**

1240

**Description**

Applies the `_summary` parameter to the Resource that is in the response (single resource or bundle).

**Name**

History

**Configuration**

`Vonk.Core.Operations.History.HistoryConfiguration`

**License token**

<http://fire.ly/vonk/plugins/history>

**Order**

4610

**Description**

Implements `_history` on system, type and instance level.

**Options**

BundleOptions, see [Search size](#)

**Name**

Version Read

**Configuration**

`Vonk.Core.Operations.History.VersionReadConfiguration`

**License token**

<http://fire.ly/vonk/plugins/history>

**Order**

4620

**Description**

Implements reading a specific version of a resource (`<base>/Patient/123/_history/v3`).

**Name**

Capability

**Configuration**

`Vonk.Core.Operations.Capability.CapabilityConfiguration`

**License token**

<http://fire.ly/vonk/plugins/capability>

**Order**

4120

**Description**

Provides the CapabilityStatement on the `<base>/metadata` endpoint. The CapabilityStatement is tailored to the FHIR version of the request. The CapabilityStatement is built dynamically by visiting all the registered implementations of `ICapabilityStatementContributor`, see [Capabilities](#).

**Name**

Conditional Create

**Configuration**

`Vonk.Core.Operations.ConditionalCrud.ConditionalCreateConfiguration`

**License token**

<http://fire.ly/vonk/plugins/conditionalcreate>

**Order**

4510

**Description**

Implements FHIR conditional create.

**Name**

Conditional Update

**Configuration**

`Vonk.Core.Operations.ConditionalCrud.ConditionalUpdateConfiguration`

**License token**

<http://fire.ly/vonk/plugins/conditionalupdate>

**Order**

4520

**Description**

Implements FHIR conditional update.

**Name**

Conditional Delete

**Configuration**

`Vonk.Core.Operations.ConditionalCrud.ConditionalDeleteConfiguration`

**License token**

<http://fire.ly/vonk/plugins/conditionaldelete>

**Order**

4530

**Description**

Implements FHIR conditional delete.

**Options**

`FhirCapabilities.ConditionalDeleteOptions`, see *FHIR Capabilities*

**Name**

Validation

**Configuration**

`Vonk.Core.Operations.Validation.ValidationConfiguration`

**License token**

<http://fire.ly/vonk/plugins/validation>

**Order**

4000

**Description**

Implements *FHIR \$validate* on type and instance level for POST: POST <base>/Patient/\$validate or POST <base>/Patient/123/\$validate.

**Name**

Instance Validation

**Configuration**

`Vonk.Core.Operations.Validation.InstanceValidationConfiguration`

**License token**

<http://fire.ly/vonk/plugins/validation>

**Order**

4840

**Description**

Implements [FHIR \\$validate](#) on instance level for GET: GET <base>/Patient/123/\$validate

**Name**

Structural Validation

**Configuration**

Vonk.Core.Operations.Validation.StructuralValidationConfiguration

**License token**

<http://fire.ly/vonk/plugins/validation>

**Order**

1227

**Description**

Validates the structure of resources sent to Firely Server (is it valid FHIR JSON or XML?).

**Name**

Prevalidation

**Configuration**

Vonk.Core.Operations.Validation.PreValidationConfiguration

**License token**

<http://fire.ly/vonk/plugins/validation>

**Order**

1228

**Description**

Validates resources sent to Firely Server against their stated profile compliance (in Resource.meta.profile). The strictness of the validation is controlled by the options.

**Options**

Validation, see [Validation](#)

**Name**

Profile filter

**Configuration**

Vonk.Core.Operations.Validation.ProfileFilterConfiguration

**License token**

<http://fire.ly/vonk/plugins/validation>

**Order**

4310

**Description**

Blocks resources that do not conform to a list of profiles.

**Options**

Validation.AllowedProfiles, see [Validation](#)

**Name**

Meta

**Configuration**

Vonk.Core.Operations.MetaOperation.MetaConfiguration

**License token**

<http://fire.ly/vonk/plugins/meta>

**Order**

5180

**Description**

Implements FHIR \$meta on instance level.

**Name**

Meta Add

**Configuration**

Vonk.Core.Operations.MetaOperation.MetaAddConfiguration

**License token**<http://fire.ly/vonk/plugins/meta>**Order**

5190

**Description**

Implements FHIR \$meta-add on instance level.

**Name**

Meta Delete

**Configuration**

Vonk.Core.Operations.MetaOperation.MetaDeleteConfiguration

**License token**<http://fire.ly/vonk/plugins/meta>**Order**

5195

**Description**

Implements FHIR \$meta-delete on instance level.

**Name**

Snapshot Generation

**Configuration**

Vonk.Core.Operations.SnapshotGeneration.SnapshotGenerationConfiguration

**License token**<http://fire.ly/vonk/plugins/snapshotgeneration>**Order**

4850

**Description**Implements [FHIR \\$snapshot](#) on a type level: POST <base>/administration/StructureDefinition/\$snapshot.**Name**

Batch

**Configuration**

Vonk.Core.Operations.Transaction.FhirBatchConfiguration

**License token**<http://fire.ly/vonk/plugins/batch>**Order**

3110

**Description**

Processes a batch Bundle by sending each entry through the rest of the processing pipeline and gathering the results.

**Options**

SizeLimits, see *Protect against large input*

**Name**

Transaction

**Configuration**

Vonk.Core.Operations.Transaction.FhirTransactionConfiguration

**License token**

<http://fire.ly/vonk/plugins/transaction>

**Order**

3120

**Description**

Process a transaction Bundle by sending each entry through the rest of the processing pipeline and gathering the results. Different from Batch, Transaction succeeds or fails as a whole. Transaction requires an implementation of Vonk.Core.Repository.IRepoTransactionService for transaction support by the underlying repository.

**Options**

- SizeLimits, see *Validation*
- Repository, see *Repository*

**Name**

LastN

**Configuration**

Vonk.Plugin.LastN.LastNConfiguration

**License token**

<http://fire.ly/vonk/plugins/lastn>

**Order**

5007

**Description**

Implements FHIR *\$lastn* on Observation resources.

**Name**

Erase

**Configuration**

Vonk.Plugin.EraseOperation.EraseOperationConfiguration

**License token**

<http://fire.ly/vonk/plugins/erase>

**Order**

5300

**Description**

Provides functionality to hard-delete FHIR resources in Firely Server database as opposed to the soft-delete used by default.

### 11.3.4 Terminology

**Name**

CodeSystem Lookup

**Configuration**

Vonk.Plugins.Terminology.[R3|R4|R5].CodeSystemLookupConfiguration

**License token**

<http://fire.ly/vonk/plugins/terminology>

**Order**

5110

**Description**

Implements FHIR [\\$lookup](#) on type level requests: POST <base>/administration/CodeSystem/\$lookup or GET <base>/administration/CodeSystem/\$lookup?...

**Name**

CodeSystem FindMatches / Compose

**Configuration**

Vonk.Plugins.Terminology.CodeSystemFindMatchesConfiguration

**License token**

<http://fire.ly/vonk/plugins/terminology>

**Order**

5220

**Description**

Implements FHIR [\\$compose](#) on type level requests: POST <base>/administration/CodeSystem/\$find-matches and on instance level requests: ``POST <base>/administration/CodeSystem/[id]/\$find-matches or GET <base>/administration/CodeSystem/[id]/\$find-matches?...

**Name**

ValueSet Validate Code

**Configuration**

Vonk.Plugins.Terminology.ValueSetValidateCodeConfiguration

**License token**

<http://fire.ly/vonk/plugins/terminology>

**Order**

5120

**Description**

Implements FHIR [\\$validate-code](#) on type level requests: POST <base>/administration/ValueSet/\$validate-code and instance level requests: GET <base>/administration/ValueSet/[id]/\$validate-code?... and POST <base>/administration/ValueSet/[id]/\$validate-code

**Name**

ValueSet Expand

**Configuration**

Vonk.Plugins.Terminology.ValueSetExpandConfiguration

**License token**

<http://fire.ly/vonk/plugins/terminology>

**Order**

5140

**Description**

Implements FHIR [\\$expand](#) on instance level requests: GET <base>/administration/ValueSet/[id]/\$expand?... and POST <base>/administration/ValueSet/[id]/\$expand and on type level requests: POST <base>/administration/ValueSet/\$expand.

**Name**

ConceptMap Translate

**Configuration**

Vonk.Plugins.Terminology.ConceptMapTranslateConfiguration

**License token**<http://fire.ly/vonk/plugins/terminology>**Order**

5260

**Description**

Implements FHIR [\\$translate](#) on instance level requests: GET <base>/administration/ConceptMap/[id]/\$translate?... and POST <base>/administration/ValueSet/[id]/\$translate and on type level requests: POST <base>/administration/ConceptMap/\$translate.

**Name**

CodeSystem Subsumes

**Configuration**

Vonk.Plugins.Terminology.CodeSystemSubsumesConfiguration

**License token**<http://fire.ly/vonk/plugins/terminology>**Order**

5280

**Description**

Implements FHIR [\\$subsumes](#) on instance level requests: GET <base>/administration/CodeSystem/[id]/\$subsumes?... and on type level requests: POST <base>/administration/CodeSystem/\$subsumes or GET <base>/administration/CodeSystem/\$subsumes?...

**Name**

CodeSystem Closure

**Configuration**

Vonk.Plugins.Terminology.CodeSystemClosureConfiguration

**License token**<http://fire.ly/vonk/plugins/terminology>**Order**

5300

**Description**

Implements FHIR [\\$closure](#) on system level requests: POST <base>/administration/\$closure



### 11.3.5 SMART on FHIR

**Name**

SMART on FHIR

**Configuration**

Vonk.Smart.SmartConfiguration.SmartConfiguration

**License token**

<http://fire.ly/vonk/plugins/smartonfhir>

**Order**

2000

**Description**

Implements SMART on FHIR authentication and authorization, see *Access control and SMART on FHIR*.

### 11.3.6 Subscriptions

**Name**

Subscriptions

**Configuration**

Vonk.Subscriptions.SubscriptionConfiguration.SubscriptionConfiguration

**License token**

<http://fire.ly/vonk/plugins/subscriptions>

**Order**

3200

**Description**

Implements the FHIR Subscriptions framework, see *Subscriptions*.

### 11.3.7 Auditing

**Name**

Username log

**Configuration**

Vonk.Plugin.Audit.UsernameLoggingConfiguration

**License token**

<http://fire.ly/vonk/plugins/audit>

**Order**

2010

**Description**

Makes the user id and name from the JWT token (if present) available for logging. See *Auditing* for more info.

**Name**

Audit logging for transactions

**Configuration**

Vonk.Plugin.Audit.AuditTransactionConfiguration

**License token**

<http://fire.ly/vonk/plugins/audit>

**Order**

3100

**Description**

Logs requests and responses for transactions to a file. See [Auditing](#) for more info.

**Name**

Audit log

**Configuration**

Vonk.Plugin.Audit.AuditConfiguration

**License token**

<http://fire.ly/vonk/plugins/audit>

**Order**

3150

**Description**

Logs requests and responses to a file. See [Auditing](#) for more info.

**Name**

AuditEvent logging for transactions

**Configuration**

Vonk.Plugin.Audit.AuditEventTransactionConfiguration

**License token**

<http://fire.ly/vonk/plugins/audit>

**Order**

3160

**Description**

Logs requests and responses for transactions to the database. See [Auditing](#) for more info.

**Name**

AuditEvent logging

**Configuration**

Vonk.Plugin.Audit.AuditEventConfiguration

**License token**

<http://fire.ly/vonk/plugins/audit>

**Order**

3170

**Description**

Logs requests and responses to the database. See [Auditing](#) for more info.

### 11.3.8 Demo UI

**Name**

Demo UI

**Configuration**

Vonk.UI.Demo.DemoUIConfiguration.DemoUIConfiguration

**License token**

<http://fire.ly/vonk/plugins/demoui>

**Order**

800

**Description**

Provides the landing page that you see when you request the base url from a browser. If you want to provide your own landing page, replace this plugin with your own. There is an example of that, see *Firely Server Plugin example - Create a new landing page*.

### 11.3.9 Documents

**Name**

Document generation

**Configuration**

Vonk.Plugins.DocumentOperation.DocumentOperationConfiguration

**License token**

<http://fire.ly/vonk/plugins/document>

**Order**

4900

**Description**

Implements FHIR `$document` : POST `<base>/Composition/$document` or GET `<base>/Composition/[id]/$document`

**Code**

[GitHub](#)

**Name**

Document signing

**Configuration**

Vonk.Plugins.SignatureService.SignatureConfiguration

**License token**

<http://fire.ly/vonk/plugins/signature>

**Order**

4899

**Description**

Signs a document generated by *\$document*.

**Name**

Document endpoint

**Configuration**

Vonk.Plugins.DocumentHandling.DocumentHandlingConfiguration

**License token**

<http://fire.ly/vonk/plugins/documenthandling>

**Order**

4950

**Description**

Allows [FHIR document bundles](#) to be posted to the base endpoint. Consult the [documentation](#) for more information.

**Name**

US Core Fetch DocumentReference operation

**Configuration**

Vonk.Plugin.DocRefOperation.DocRefOperationConfiguration

**License token**

<http://fire.ly/vonk/plugins/docref>

**Order**

5350

**Description**

Allows fetching DocumentReference resources related to a patient. Implements the [US Core Fetch DocumentReference](#) operation. Generating a document is not yet supported.

### 11.3.10 Conversion

**Name**

Format conversion

**Configuration**

Vonk.Plugins.ConvertOperation.ConvertOperationConfiguration

**License token**

<http://fire.ly/vonk/plugins/convert>

**Order**

4600

**Description**

Implements FHIR [\\$convert](#) : POST <base>/\$convert to convert between JSON and XML representation.

### 11.3.11 Binary

**Name**

Binary wrapper (Encode)

**Configuration**

Vonk.Plugins.BinaryWrapper.BinaryEncodeConfiguration

**License token**

<http://fire.ly/vonk/plugins/binarywrapper>

**Order**

1112

**Description**

Wraps an incoming binary format in a Binary resource for further processing by the pipeline.

**Settings**

```
"Vonk.Plugin.BinaryWrapper":{
  "RestrictToMimeTypes": ["application/pdf", "text/plain", "image/png",
↪ "image/jpeg"]
},
```

**Name**

Binary wrapper (Decode)

**Configuration**

Vonk.Plugins.BinaryWrapper.BinaryDecodeConfiguration

**License token**

<http://fire.ly/vonk/plugins/binarywrapper>

**Order**

1122

**Description**

Implements GET <base>/Binary/<id>, retrieve back the Binary resource in its native format.

### 11.3.12 Repository implementations

**Name**

Memory Repository

**Configuration**

Vonk.Repository.MemoryConfiguration

**license token**

<http://fire.ly/vonk/plugins/repository/memory>

**Order**

210

**Description**

Implements a repository in working memory that fully supports all of the capabilities of Firely Server. This implementation is mainly used for unittesting.

**Name**

Memory Administration Repository

**Configuration**

Vonk.Repository.MemoryAdministrationConfiguration

**license token**

<http://fire.ly/vonk/plugins/repository/memory>

**Order**

211

**Description**

Implements a repository in working memory for the Administration API. This implementation is mainly used for unittesting.

**Name**

MongoDb Repository

**Configuration**

`Vonk.Repository.MongoDbConfiguration`

**license token**

<http://fire.ly/vonk/plugins/repository/mongo-db>

**Order**

230

**Description**

Implements a repository in MongoDB that fully supports all of the capabilities of Firely Server.

**Name**

MongoDb Administration Repository

**Configuration**

`Vonk.Repository.MemoryAdministrationConfiguration`

**license token**

<http://fire.ly/vonk/plugins/repository/mongo-db>

**Order**

231

**Description**

Implements a repository in MongoDB for the Administration API.

**Name**

SQLite Repository

**Configuration**

`Vonk.Repository.SqliteConfiguration`

**license token**

<http://fire.ly/vonk/plugins/repository/sqlite>

**Order**

240

**Description**

Implements a repository in SQLite that fully supports all of the capabilities of Firely Server.

**Name**

SQLite Administration Repository

**Configuration**

`Vonk.Repository.SqliteAdministrationConfiguration`

**license token**

<http://fire.ly/vonk/plugins/repository/sqlite>

**Order**

241

**Description**

Implements a repository in SQLite for the Administration API.

**Name**

SQL Server Repository

**Configuration**

`Vonk.Repository.SqlConfiguration`

**license token**<http://fire.ly/vonk/plugins/repository/sql-server>**Order**

220

**Description**

Implements a repository in SQL Server that fully supports all of the capabilities of Firely Server.

**Name**

SQL Server Administration Repository

**Configuration**

Vonk.Repository.SqlAdministrationConfiguration

**license token**<http://fire.ly/vonk/plugins/repository/sql-server>**Order**

221

**Description**

Implements a repository in SQL Server for the Administration API.

### 11.3.13 Administration API

**Name**

Administration API

**Configuration**

Vonk.Administration.Api.AdministrationOperationConfiguration

**license token**<http://fire.ly/vonk/plugins/administration>**Order**

1160

**Description**

Sets up a sequence of plugins for the Administration API. Administration API is different from general plugins since it branches off of the regular processing pipeline and sets up a second pipeline for the /administration endpoint.

**Name**

Fhir STU3 Administration services

**Configuration**

Vonk.Administration.FhirR3.RepositoryConfigurationR3

**license token**<http://fire.ly/vonk/plugins/administration/fhrr3>**Order**

4310

**Description**

Implements support services to work with FHIR STU3 conformance resources in the Administration API.

**Name**

Fhir R4 Administration services

**Configuration**

Vonk.Administration.FhirR4.RepositoryConfigurationR4

**license token**

<http://fire.ly/vonk/plugins/administration/fhirr4>

**Order**

4310

**Description**

Implements support services to work with FHIR R4 conformance resources in the Administration API.

### 11.3.14 Bulk Data

**Name**

System Bulk Data Export

**Configuration**

Vonk.Plugin.BulkDataExport.SystemBulkDataExportConfiguration

**license token**

<http://fire.ly/vonk/plugins/bulk-data-export>

**Order**

5003

**Description**

Support for system-level \$export operation. See *Bulk Data Export*.

**Name**

Group Bulk Data Export

**Configuration**

Vonk.Plugin.BulkDataExport.GroupBulkDataExportConfiguration

**license token**

<http://fire.ly/vonk/plugins/bulk-data-export>

**Order**

5004

**Description**

Support for instance-level \$export operation. See *Bulk Data Export*.

**Name**

Patient Bulk Data Export

**Configuration**

Vonk.Plugin.BulkDataExport.PatientBulkDataExportConfiguration

**license token**

<http://fire.ly/vonk/plugins/bulk-data-export>

**Order**

5005

**Description**

Support for type-level \$export operation. See *Bulk Data Export*.

**Name**

Patient everything



**Configuration**

Vonk.Plugin.PatientEverything

**license token**

<http://fire.ly/vonk/plugins/patient-everything>

**Order**

5006

**Description**

Request a Patient record. See *Patient \$everything*.

### 11.3.15 X-Provenance header

**Name**

X-Provenance header

**Configuration**

Vonk.Plugin.Operations.Provenance.ProvenanceHeaderConfiguration

**license token**

<http://fire.ly/vonk/plugins/transaction>

**Order**

1230

**Description**

Support for the X-Provenance header that adds a Provenance resource upon creating or updating another resource. See *X-Provenance header*.

## 11.4 Important classes and interfaces

If you want to develop a plugin for Firely Server, there are a couple of classes that you will probably interact with. This page lists those classes, with an explanation of each.

### 11.4.1 IResource

**namespace**

Vonk.Core.Common

**purpose**

IResource is the abstraction for all FHIR resources in Firely Server. It is used in the request and the response, and thereby all through the pipeline. It allows you to program against resources in different Fhir.NET API (the Resource class is defined in each version separately), as well as against resources that do not even have a POCO implementation.

```
/// <summary>
/// Abstraction of a resource in some format. Specifies the properties of a Resource
↳ that Firely Server needs to read and maintain.
/// <para>Future: may be extended with Tags and Labels.</para>
/// </summary>
public interface IResource : ISourceNode
{
    /// <summary>
```

(continues on next page)

```

    /// Type of resource, e.g. Patient or AllergyIntolerance.
    /// </summary>
    string Type { get; }

    /// <summary>
    /// Logical identity of resource, e.g. 'example' or 'e3f5b0b8-4570-4e4c-b597-
    ↪e6523aff3a19'. Does not contain the resourcetype.
    /// Refers to Resource.id
    /// IResource is immutable, so to update this, use resourceWithNewId = this.SetId(),
    ↪ from IResourceExtensions.
    /// In the context of a repository, consider IResourceChangeRepository.EnsureMeta().
    /// </summary>
    string Id { get; }

    /// <summary>
    /// Version of resource. Refers to Resource.meta.versionId.
    /// IResource is immutable, so to update this, use resourceWithNewVersion = this.
    ↪SetVersion(), from IResourceExtensions.
    /// In the context of a repository, consider IResourceChangeRepository.EnsureMeta().
    /// </summary>
    string Version { get; }

    /// <summary>
    /// Model that the resource was defined in.
    /// Common models are the different versions of FHIR, defined in <see cref=
    ↪"VonkConstants.Model"/>
    /// </summary>
    string InformationModel { get; }

    /// <summary>
    /// When was the resource last updated?
    /// Refers to Resource.meta.lastUpdated.
    /// IResource is immutable, so to update this, use resourceWithNewLastUpdated =
    ↪this.SetLastUpdated(DateTimeOffset) from IResourceExtensions.
    /// In the context of a repository, consider IResourceChangeRepository.EnsureMeta().
    /// </summary>
    DateTimeOffset? LastUpdated { get; }

    /// <summary>
    /// Is this a contained resource, or a container resource?
    /// A resource is a container resource if it is not contained. Even if it has no
    ↪contained resources embedded.
    /// </summary>
    ResourceContained Contained { get; }

    /// <summary>
    /// Direct access to contained resources, if any. Prefer to return an empty list
    ↪otherwise.
    /// Refers to DomainResource.contained.
    /// </summary>
    IEnumerable<IResource> ContainedResources { get; }
}

```

If you work with a POCO, you can use the extension method `ToIResource()` from `Vonk.Fhir.R3` to adapt it to an `IResource`:

```
var patientPoco = new Patient(); //Requires Hl7.Fhir.Model
var resource = patientPoco.ToIResource();
```

`IResource` is immutable, so changes will always result in a new instance. Changes can usually be applied with extension methods on `ISourceNode`, found in `Vonk.Core.ElementModel.ISourceNodeExtensions`. There are also several extension methods specifically for `IResource` in `Vonk.Core.Common.IResourceExtensions`:

```
var updatedResource = oldResource.Add(SourceNode.Valued("someElement", "someValue");
//Continue with updatedResource, since oldResource will not have the change.
```

## IResource extension methods

`IResource` has a whole list of extension methods for manipulating them, conversion between `ISourceNode` and `IResource` and caching objects within it. All these methods are in the namespace `Vonk.Core.Common.IResourceExtensions`. Please check the `///`-comments on the methods for more information.

### 11.4.2 ElementModel, manipulating IResource and ISourceNode

As stated in [IResource](#), `IResource` is immutable, and so is `ISourceNode` that it is based on. Still, there are cases where you would want to manipulate a resource, e.g. add, change or remove elements. The namespace `Vonk.Core.ElementModel` has methods to do so. All of these methods do *NOT* change the original structure (the input `IResource` or `ISourceNode`), but instead *return* an updated structure.

## ISourceNode manipulation

All the `ISourceNode` extension methods can be used on `IResource` as well. If you need an `IResource` as result, just turn the resulting `ISourceNode` to an `IResource` again. So if you added an element to an `ISourceNode`:

```
IResource result = input.AddIfNotExists(SourceNode.Valued("active", "false")).
    ↪ ToIResource(input.InformationModel);
```

For some of the more common extension methods we provide an overload on `IResource` that does this for you, like `IResource.Patch(...)`

All the methods below are in the namespace `Vonk.Core.ElementModel.ISourceNodeExtensions`:

**Add**(this *ISourceNode original*, *ISourceNode newChild*) → *ISourceNode*

Add the *newChild* as a child node to the *original*. It will be added at the end of the Children.

**Add**(this *ISourceNode original*, *ITypedElement newChild*) → *ISourceNode*

Overload of `Add(ISourceNode newChild)` that lets you add an `ITypedElement` as new child.

**AddIf**(this *ISourceNode original*, *ISourceNode newChild*, *Func<ISourceNode, bool> addIf*) → *ISourceNode*

Add the *newChild* as a child node to the *original* if the *addIf* predicate on *original* is met. It will be added at the end of the Children.

**Add**(this *ISourceNode original*, *TypedElement newChild*, *Func<ISourceNode, bool> addIf*) → *ISourceNode*

Similar to `AddIf(ISourceNode newChild, Func<ISourceNode, bool> addIf)` that lets you add an `ITypedElement` as new child.

**AddIfNotExists**(*this ISourceNode original, ISourceNode newChild*) → ISourceNode

Add the *newChild* as a child node to the *original* if there is no child with the same name yet. It will be added at the end of the Children.

**AddIfNotExists**(*this ISourceNode original, ISourceNode newChild, Func<ISourceNode, bool> exists*) → ISourceNode

Add the *newChild* as a child node to the *original* if the *exists* predicate on *original* is not satisfied. This is like `AddIfNotExist(ISourceNode newChild)`, but here you get to specify what 'exists' means. It will be added at the end of the Children.

**AddIfNotExists**(*this ISourceNode original, string location, ISourceNode newChild*) → ISourceNode

Navigate to *location*. Then add the *newChild* as a child node to the *original* if there is no child with the same name yet.

**AddIfNotExists**(*this ISourceNode original, ISourceNode newChild, string location, Func<ISourceNode, bool> exists*) → ISourceNode

Navigate to *location*. Then add the *newChild* as a child node if the *exists* predicate on the current node is not satisfied.

**AnnotateWith<T>**(*this ISourceNode original, T annotation, bool hideExisting = false*) → ISourceNode

Add an annotation of type *T* to the *original*. When *hideExisting* == true, any existing annotations of type *T* are not visible anymore on the returned ISourceNode.

**GetBoundAnnotation<T>**(*this ISourceNode original, ) where T : class, IBoundAnnotation* → T

Retrieve an annotation that is bound directly to *original*, not to any of the nodes it may decorate. (ISourceNode is immutable, to changes are usually a pile of wrappers around the original SourceNode, and each of the wrappers can add / replace annotations.)

**RemoveEmptyNodes**(*this ISourceNode original, ISourceNode newChild*) → ISourceNode

Remove any nodes that have no value or children. This happens recursively: if a node has only children with empty values, it will be removed as well. This way the returned ISourceNode conforms to the invariant in the FHIR specification that an element either has a value or children.

**RemoveEmptyNodes**(*this ISourceNode original, ISourceNode newChild, string location*) → ISourceNode

Remove any nodes that have no value or children, from the specified *location* downwards. This happens recursively: if a node has only children with empty values, it will be removed as well.

**Child**(*this ISourceNode original, string name, int arrayIndex = 0*) → ISourceNode

Convenience method to get the child with name *name* at position *arrayIndex*. Usually used to get a child of which you know there is only one: `patientNode.Child("active")`

**ChildString**(*this ISourceNode original, string name, int arrayIndex = 0*) → ISourceNode

Convenience method to get the value of the child with name *name* at position *arrayIndex*. Usually used to get a child of which you know there is only one: `patientNode.ChildString("id")`

**ForceAdd**(*this ISourceNode original, string addAt, ISourceNode newChild*) → ISourceNode

Add the *newChild* at location *addAt*. Create the intermediate nodes if necessary.

**AddOrReplace**(*this ISourceNode original, Func<ISourceNode, bool> match, ISourceNode toAdd, Func<ISourceNode, ISourceNode> replace*) → ISourceNode

Find any child nodes of *original* that match the *match* predicate. Apply *replace* to them. If none are found, add *toAdd* as new child.

**AddOrReplace**(*this ISourceNode original, ISourceNode toAdd, Func<ISourceNode, ISourceNode> replace*) → ISourceNode

Optimized overload of the previous method for matching on the node name. It will perform **replace** on any child node of **original** with the same name as **toAdd**. If none are found it will add **toAdd** as new child node.

**Remove**(*this ISourceNode original, string location*) → ISourceNode

Remove the node at **location**, if any. If that results in parent nodes becoming empty (no Text, no Children), those are removed as well.

**SelectNodes**(*this ISourceNode original, string fhirPath*) → IEnumerable<ISourceNode>

Run **fhirPath** over the **original**, but with the limitations of untyped nodes. It will return the matching nodes. Use **valueDateTime/valueBoolean** instead of just 'value' for choice types. Only use this method if you are familiar with the differences in the naming of nodes between **ISourceNode** and **ITypedElement**.

**SelectText**(*this ISourceNode original, string fhirPath*) → string

Run **fhirPath** over the **original**, but with the limitations of untyped nodes. Returns the **Text** of the first matching node. Use **valueDateTime/valueBoolean** instead of just 'value' for choice types. Only use this method if you are familiar with the differences in the naming of nodes between **ISourceNode** and **ITypedElement**.

**Patch**(*this ISourceNode original, string location, Func<ISourceNode, ISourceNode> patch*) → ISourceNode

Find any nodes at **location** and apply **patch** to them. For **patch** you can use other methods listed here like **Rename**, **Add** or **Revalue**. **location** is evaluated as a fhirpath statement, with the limitations of untyped nodes.

**Patch**(*this ISourceNode original, string[] locations, Func<ISourceNode, ISourceNode> patch*) → ISourceNode

Find any nodes having one of the **locations** as their **Location** and apply **patch** to them. If you don't know exact locations, use **original.Patch(location, patch)**, see above.

**ForcePatch**(*this ISourceNode original, string forcePath, Func<ISourceNode, ISourceNode> patch*) → ISourceNode

Enforce that **forcePath** exists. Then patch the resulting node(s) with **patch**.

**ForcePatchAt**(*this ISourceNode original, string fromLocation, string forcePath, Func<ISourceNode, ISourceNode> patch*) → ISourceNode

For each node matching the **fromLocation**: enforce that **fromLocation.forcePath** exists, then patch the resulting node(s) with **patch**. E.g. **someBundle.ForcePatchAt("entry", "request", node => node.Add(SourceNode.Valued("url", "someUrl"))** will add **request.url** with value "someUrl" to every entry.

**Relocate**(*this ISourceNode original, string newLocation*) → ISourceNode

Set **original.Location** to the **newLocation**, and update all its descendants' **Location** properties recursively.

**Rename**(*this ISourceNode original, string newName*) → ISourceNode

Set **original.Name** to the **newName**.

**Revalue**(*this ISourceNode original, string newValue*) → ISourceNode

Set **original.Text** to **newValue**.

**Revalue**(*this ISourceNode original, Dictionary<string, string> replacements*) → ISourceNode

**replacements** is a dictionary of location + **newValue**. On each matching location under **original**, the value will be set to the according **newValue** from **replacements**.

**AnnotateWithSourceNode**(*this ISourceNode original*) → ISourceNode

Add **original** as annotation to itself. Very specific use case.

## ITypedElement manipulation

All the methods below are in the namespace `Vonk.Core.ElementModel.ITypedElementExtensions`:

**Add**(*this ITypedElement original, ITypedElement newChild, Func<ITypedElement, bool> addIf*) → `ISourceNode`

Add `newChild` as child to `original` if `addIf` on `original` evaluates to true. Convenience overload of `ISourceNodeExtensions.Add(ISourceNode, ITypedElement, Func<ISourceNode, bool>)`

**Add**(*this ITypedElement original, ITypedElement newChild*) → `ISourceNode`

Add `newChild` as child to `original`. Convenience overload of `ISourceNodeExtensions.Add(ISourceNode, ITypedElement)`

**AddIfNotExists**(*this ITypedElement original, ITypedElement newChild*) → `ISourceNode`

Add `newChild` as child to `original` if no child with the same name exists yet. Convenience overload of `ISourceNodeExtensions.AddIfNotExists(ISourceNode, ITypedElement)`

**AddIf**(*this ITypedElement original, ISourceNode newChild, Func<ITypedElement, bool> addIf*) → `ISourceNode`

Add `newChild` as child to `original` if `addIf` on `original` evaluates to true. Convenience overload of `ISourceNodeExtensions.AddIf(ISourceNode, ISourceNode, Func<ISourceNode, bool>)`

### method

`Add(this ITypedElement original, ISourceNode newChild)` Add `newChild` as child to `original`.

### method

`AddIfNotExists(this ITypedElement original, ISourceNode newChild)` Add `newChild` as child to `original` if no child with the same name exists yet. Convenience overload of `AddIfNotExists(ITypedElement, ITypedElement)`

**Cache**(*this ITypedElement original*) → `ITypedElement`

Prevent recalculation of the Children upon every access.

**Child**(*this ITypedElement element, string name, int arrayIndex = 0*) → `ITypedElement`

Returns n-th child with the specified name, if any.

**ChildString**(*this ITypedElement element, string name, int arrayIndex = 0*) → `string`

Returns the value of the n-th child with the specified name as string, if any.

**DefinitionSummary**(*this ITypedElement element, IStructureDefinitionSummaryProvider provider*) → `IStructureDefinitionSummary`

Returns the summary for the actual type of the element. Especially useful if the element is of a choicetype.

**AddParent**(*this ITypedElement element*) → `ITypedElement`

Add `Vonk.Core.ElementModel.IParentProvider` annotations to `element` and its descendants.

**GetParent**(*this ITypedElement element*) → `ITypedElement`

Get the parent of this element, through the `Vonk.Core.ElementModel.IParentProvider` annotation (if present).

**AddTreePath**(*this ITypedElement element*) → `ITypedElement`

Add the `Vonk.Core.ElementModel.ITreePathGenerator` annotation. `TreePath` is the Location without any indexes (no [n] at the end).

**GetTreePath**(*this ITypedElement element*) → `string`

Get the value of the `Vonk.Core.ElementModel.ITreePathGenerator` annotation, if present. `TreePath` is the Location without any indexes (no [n] at the end).

### 11.4.3 IVonkContext

#### namespace

Vonk.Core.Context

#### purpose

IVonkContext is the Vonk-specific counterpart to HttpContext from ASP.NET Core. It contains an IVonkRequest and IVonkResponse object that allow you to get information from the request and set results in the response, both in Firely Server terms.

Have IVonkContext injected in the method where you need it. Use a *configuration class* to call this method from the pipeline and have the actual context injected. A more complete template is found at *Template for a plugin*.

```
public class SomeService
{
    public async Task DoMyOperation(IVonkContext vonkContext)
    {
        //...
    }
}

public static class SomeServiceConfiguration
{
    public static IApplicationBuilder UseMyOperation(this IApplicationBuilder app)
    {
        return app.UseVonkInteractionAsync<SomeService>((svc, context) => svc.
        DoMyOperation(context));
    }
}
```

If you also need access to the raw HttpContext, you have two options:

1. The IVonkContext.HttpContext() extension method gives you the original HttpContext. Be aware though:
  - Situations may arise where there is no HttpContext, so be prepared for that.
  - When handling batches or transactions, an IVonkContext is created for each entry in the bundle. But they all refer to the same original HttpContext.
2. Create a normal ASP.NET Core Middleware class and access the IVonkContext with the extension method Vonk() on HttpRequest. A more complete template is found at *Returning non-FHIR content from a plugin*.

```
public class SomeMiddleware
{
    public SomeMiddleware(RequestDelegate next)
    {
        //...
    }

    public async Task Invoke(HttpContext httpContext)
    {
        var vonkContext = httpContext.Vonk();
        //...
    }
}
```

(continues on next page)

(continued from previous page)

```

public static class SomeMiddlewareConfiguration
{
    public static IApplicationBuilder UseSomeMiddleware(this IApplicationBuilder app)
    {
        return app.UseMiddleware<SomeMiddleware>(); //Just plain ASP.NET Core, ↵
        ↵ nothing Firely Server specific here.
    }
}

```

IVonkContext has three major parts, that are explained below. The InformationModel tells you the FHIR version for which the request was made.

```

public interface IVonkContext
{
    IVonkRequest Request {get;}

    IArgumentCollection Arguments {get;}

    IVonkResponse Response {get;}

    string InformationModel {get;}
}

```

And because you frequently need the parts instead of the context itself, there is an extension method on IVonkContext:

```

public (IVonkRequest request, IArgumentCollection args, IVonkResponse respons) ↵
    ↵ Parts(this IVonkContext vonkContext)

```

## IVonkRequest

### namespace

Vonk.Core.Context

### purpose

Get information about the request made, in Firely Server / FHIR terms.

You can access the current IVonkRequest through the *IVonkContext*. Its properties are:

```

public interface IVonkRequest
{
    string Path { get; }
    string Method { get; }
    string CustomOperation { get; }
    VonkInteraction Interaction { get; }
    RequestPayload Payload { get; set; }
}

```

Path and Method relate directly to the equivalents on HttpContext. Interaction tells you which of the FHIR RESTful interactions was called. CustomOperation is only filled if one of the custom operations was invoked, like e.g. \$validate. All of these can be filtered by the *InteractionHandlerAttribute*, so you typically don't need to inspect them manually.

Payload indirectly contains the resource that was sent in the body of the request. You are advised to only use the extension methods to access it:



```
public static bool TryGetPayload(this IVonkRequest request, out IResource resource)
```

TryGetPayload is useful if your code wants to act on the payload *if it is present*, but does not care if it is not.

```
public void ThisMethodActsOnThePayloadIfPresent(IVonkContext vonkContext)
{
    var (request, args, response) = vonkContext.Parts();
    if (request.TryGetPayload(response, out var resource))
    {
        // do something with the resource.
    }
}
```

```
public static bool GetRequiredPayload(this IVonkRequest request, IVonkResponse response,
    ↪out IResource resource)
```

GetRequiredPayload is useful if your code expects the payload to be present. It will set the appropriate response code and OperationOutcome on the provided response if it is not present or could not be parsed. Then you can choose to end the pipeline and thus return the error to the user.

```
public void ThisMethodNeedsAPayload(IVonkContext vonkContext)
{
    var (request, args, response) = vonkContext.Parts();
    if (!request.GetRequiredPayload(response, out var resource))
    {
        return; //If you return with an error code in the response, Firely Server will end
    ↪the pipeline
    }
    // do something with the resource.
}
```

If you want to **change** the payload, assign a whole new one. Generally you would want to change something to the old payload. But IResource is immutable, so changes to it yield a new instance. That leads to this pattern

```
if (request.TryGetPayload(response, out var resource)
{
    //Explicit typing of variables for clarity, normally you would use 'var'.
    ISourceNode updatedNode = resource.Add(SourceNode.Valued("someElement", "someValue");
    IResource updatedResource = updatedNode.ToIResource();
    request.Payload = updatedResource.ToPayload();
}
```

### IArgumentCollection, IArgument

#### namespace

Vonk.Core.Context

#### purpose

Access arguments provided in the request.

The `IVonkContext.Arguments` property contains all the arguments from the request, from the various places:

1. The path segments: `/Patient/123/_history/v1` will translate to three arguments, `_type`, `_id` and `_version`.
2. The query parameters: `?name=Fred&active=true` will translate to two arguments, `name` and `active`.
3. The headers:
  1. `If-None-Exists = identifier=abc&active=true` will translate to two arguments, `identifier` and `active`.
  2. `If-Modified-Since`, `If-None-Match`, `If-Match`: will each translate to one argument

An individual argument will tell you its name (`ArgumentName`), raw value (`ArgumentValue`) and where it came from (`Source`).

### Handling arguments

An argument by default has a `Status` of `Unhandled`.

If an argument is of interest to the operation you implement in your plugin, you can handle the argument. It is important to mark arguments handled if:

- you handled them
- or the handling is not relevant anymore because of some error you encountered

In both cases you simply set the `Status` to `Handled`.

If an argument is incorrect, you can set its status to `Error` and set the `Issue` to report to the client what the problem was. These issues will be accumulated in the response by Firely Server automatically.

Any argument that is not handled will automatically be reported as such in an `OperationOutcome`.

Useful extension methods:

```
IArgument.Handled()
IArgument.Warning(string message, Issue issue)
IArgument.Error(string message, Issue issue)
```

Firely Server has a lot of issues predefined in `Vonk.Core.Support.VonkIssues`.

### IVonkResponse

#### namespace

Vonk.Core.Context

#### purpose

Inspect response values set by other middleware, or set it yourself.

```
public interface IVonkResponse
{
    Dictionary<VonkResultHeader, string> Headers { get; }
    int HttpResult { get; set; }
    OperationOutcome Outcome { get; }
    IResource Payload { get; set; }
}
```

If your operation provides a response, you should:

1. Set the response code `HttpResult`.
2. Provide a resource in the `Payload`, if applicable.
3. Add an issue if something is wrong.

If you just listen in on the pipeline, you can check the values of the response. Besides that, the *InteractionHandlerAttribute* allows you to filter on the `HttpStatus` of the response.

## IFormatter

### namespace

Vonk.Core.Context.Format

### purpose

Serialize response resource in requested format to the body of the `HttpContext.Response`. Although this interface is public, you should never need it yourself, since the *VonkToHttp plugin* takes care of this for you.

## 11.4.4 Interaction Handling

In the configuration of a plugin you specify on which interaction(s) the plugin should act. That can be done with an attribute on the main method of the service in the plugin, or with a fluent interface on `IApplicationBuilder`.

## InteractionHandlerAttribute

### namespace

Vonk.Core.Pluggability

### purpose

Add an `[InteractionHandler]` attribute to a method to specify when the method has to be called. You specify this by providing values that the `IVonkContext` should match.

Without any arguments, the method will be called for every possible interaction.

```
[InteractionHandler()]
public async Task DoMyOperation(IVonkContext vonkContext)
```

You can specify different filters, and combine them at will:

- Specific interaction(s): `[InteractionHandler(Interaction = VonkInteraction.type_create | VonkInteraction.instance_update)]`
- Specific FHIR version(s) of the request: `[InteractionHandler(InformationModel = VonkConstants.Model.FhirR4)]`

- Specific resource type(s): `[InteractionHandler(AcceptedTypes = new["Patient", "Observation"])]`
- Specific custom operation: `[InteractionHandler(Interaction = VonkInteraction.all_custom, CustomOperation = "myCustomOperation")]`. Note that the \$ that is used on the url is not included in the name of the custom operation here.
- Specific http method: `[InteractionHandler(Method = "POST")]`
- Specific statuscode(s) on the response: `[InteractionHandler(StatusCode = new[]{200, 201})]`

Now to configure your service to be a processor in the Firely Server pipeline, you use `UseVonkInteraction[Async]()`:

```
public static class MyOperationConfiguration
{
    public static IApplicationBuilder UseMyOperation(this IApplicationBuilder app)
    {
        return app.UseVonkInteractionAsync<MyService>((svc, ctx) => svc.
        ↪ DoMyOperation(ctx));
    }
}
```

### InteractionHandler fluent interface

Because `InteractionHandler` is an attribute, you can only use constant values. If that is not what you want, you can use the fluent interface in the `configuration class` instead. The code below shows the same filters as above, although you typically would not use all of them together (e.g. the PUT excludes `type_create`).

```
public static class MyOperationConfiguration
{
    public static IApplicationBuilder UseMyOperation(this IApplicationBuilder app)
    {
        return app
            .OnInteraction(VonkInteraction.type_create | VonkInteraction.instance_update)
            .AndInformationModel(VonkConstants.Model.FhirR4)
            .AndResourceTypes(new[] { "Patient", "Observation" })
            .AndStatusCodes(new[] { 200, 201 })
            .AndMethod("PUT")
            .HandleAsyncWith<MyService>((svc, ctx) => svc.DoMyOperation(ctx));
    }
}
```

Other `Handle...` methods allow you to define a pre-handler (that checks or alters the request before the actual operation) or a post-handler (that checks or alters the response after the actual operation), either synchronously or asynchronously.

If you have a very specific filter that is not covered by these methods, you can specify it directly with a function on the `IVonkContext` that returns a boolean whether or not to call your operation.

```
app
    .On(ctx => MyVerySpecificFilter(ctx))
    .Handle...
```

**Attention:** The filter you specify is called for **every** request. So make sure you don't do any heavy calculations or I/O.

## IApplicationBuilder extension methods

**UseVonkInteraction**`<TService>(this IApplicationBuilder app, Expression<Action<<TService, IVonkContext>> handler, OperationType operationType = OperationType.Handler) -> IApplicationBuilder`

Handle the request with the handler method when the request matches the `InteractionHandler` attribute on the handler method. The `OperationType` may also specify `PreHandler` or `PostHandler`. If you need to do anything lengthy (I/O, computation), use the `Async` variant of this method.

**UseVonkInteractionAsync**`<TService>(this IApplicationBuilder app, Expression<Func<TService, IVonkContext, T.Task>> handler, OperationType operationType = OperationType.Handler) -> IApplicationBuilder`

Handle the request with the asynchronous handler method when the request matches the `InteractionHandler` attribute on the handler method. The `OperationType` may also specify `PreHandler` or `PostHandler`.

**OnInteraction**`(this IApplicationBuilder app, VonkInteraction interaction) -> VonkApplicationBuilder`

Used for fluent configuration of middleware. This is one of two methods to enter the `VonkApplicationBuilder`, see [VonkApplicationBuilder extension methods](#). It requires you to choose an interaction to act on. If you need your services to act on every interaction, choose `VonkInteraction.all`.

**OnCustomInteraction**`(this IApplicationBuilder app, VonkInteraction interaction, string custom) -> VonkApplicationBuilder`

Used for fluent configuration of middleware. This is one of two methods to enter the `VonkApplicationBuilder`, see [VonkApplicationBuilder extension methods](#). It requires you to choose an interaction to act on. This should be one of the `VonkInteraction.all_custom` interactions. `custom` is the name of the custom interaction to act on, without the preceding '\$'.

## VonkApplicationBuilder extension methods

`VonkApplicationBuilder` is used to fluently configure your middleware. It has methods to filter the requests that your middleware should respond to. Then it has a couple of `*Handle...` methods to transform your service into middleware for the pipeline, and return to the `IApplicationBuilder` interface.

**AndInteraction**`(this VonkApplicationBuilder app, VonkInteraction interaction) -> VonkApplicationBuilder`

Specify an interaction to act on.

**AndResourceTypes**`(this VonkApplicationBuilder app, params string[] resourceTypes) -> VonkApplicationBuilder`

Specify the resourcetypes to act on.

**AndStatusCodes**`(this VonkApplicationBuilder app, params int[] statusCodes) -> VonkApplicationBuilder`

Specify the statuscode(s) of the response to act on. This is mainly useful for posthandlers.

**AndMethod**`(this VonkApplicationBuilder app, string method) -> VonkApplicationBuilder`

Specify the http method (GET, PUT, etc) to act on.

**AndInformationModel**`(this VonkApplicationBuilder app, string model) -> VonkApplicationBuilder`

If your service can only act on one FHIR version, specify it with this method. Common values for `model` are `VonkConstants.Model.FhirR3` and `VonkConstants.Model.FhirR4`.

```
PreHandleAsyncWith<TService>(this VonkApplicationBuilder app, Expression<Func<TService, IVonkContext, T.Task>> preHandler) -> IApplicationBuilder
```

Mark the preHandler method as a prehandler, so it will act on the IVonkContext and send it further down the pipeline.

```
PreHandleWith<TService>(this VonkApplicationBuilder app, Expression<Action<TService, IVonkContext>> preHandler) -> IApplicationBuilder
```

Synchronous version of PreHandleAsyncWith for synchronous preHandler methods.

```
HandleAsyncWith<TService>(this VonkApplicationBuilder app, Expression<Func<TService, IVonkContext, T.Task>> handler) -> IApplicationBuilder
```

Mark the handler method as a handler, so it will act on the IVonkContext, provide a response and end the pipeline for the request.

```
HandleWith<TService>(this VonkApplicationBuilder app, Expression<Action<TService, IVonkContext>> handler)
```

Synchronous version of HandleAsyncWith for synchronous handler methods.

```
PostHandleAsyncWith<TService>(this VonkApplicationBuilder app, Expression<Func<TService, IVonkContext, T.Task>> postHandler) -> IApplicationBuilder
```

Mark the postHandler method as a posthandler, so it will pass on the IVonkContext to the rest of the pipeline, and on the way back through the pipeline inspect or modify the response. Make sure that the VonkConfiguration order you have for this is lower than whatever action you need to post-handle.

```
PostHandleWith<TService>(this VonkApplicationBuilder app, Expression<Action<TService, IVonkContext>> postHandler) -> IApplicationBuilder
```

Synchronous version of PostHandleAsyncWith for synchronous postHandler methods.

## 11.4.5 Pipeline configuration

### VonkConfigurationAttribute

#### namespace

Vonk.Core.Pluggability

#### purpose

This attribute is used on a static class to make Firely Server recognize it as part of the configuration of the processing pipeline. See *Configuration classes*.

#### properties

- Order: Determines the place in the pipeline. See *The order of plugins* for background.

## 11.4.6 Repository interfaces

### ISearchRepository

Vonk.Core.Repository.ISearchRepository is central in the functioning of Firely Server. It defines all read-access to the underlying repository, being one of Firely Server's own database implementations or a Facade implementation.

It has a single method:

```
Task<SearchResult> Search(IArgumentCollection arguments, SearchOptions options);
```

## Using ISearchRepository

1. Have it injected by the dependency injection into the class where you need it:

```
public MyService(ISearchRepository searchRepository)
{
    _searchRepository = searchRepository;
}
```

**Note:** Implementations of ISearchRepository have a *Scoped* lifetime, so MyService should also be registered as Scoped:

```
public IServiceCollection AddMyServices(this IServiceCollection services)
{
    services.TryAddScoped<MyService>();
    return services;
}
```

2. Prepare search arguments that express what you are looking for. Search arguments are effectively key-value pairs as if they came from the querystring on the request. So the key must be the code of a supported Searchparameter.

```
var args = new ArgumentCollection(new IArgument[]
{
    new Argument(ArgumentSource.Internal, ArgumentNames.resourceType, "Patient")
    ↳ {MustHandle = true}, // MustHandle = true is optional
    new Argument(ArgumentSource.Internal, "name", "Fred") {MustHandle = true} // ↳
    ↳ MustHandle = true is optional
}).AddCount(20);
```

**Note:** The Search implementation will in general update the arguments, especially their Status property and the Issue if something went wrong. So be careful with reuse of arguments. Use IArgumentCollection.Clone() if necessary.

3. Prepare search options that guide the search. Usually you can use one of the predefined options on the SearchOptions class.

```
var options = SearchOptions.Latest(vonkContext.ServerBase, vonkContext.Request.
    ↳ Interaction, vonkContext.InformationModel);
```

4. Execute the search.

```
var searchResult = await _searchRepository.Search(args, options);
```

5. Check the status of the arguments, especially if they could not be ignored (MustHandle = true). Because this is a common pattern, there is an extension method CheckHandled that throws a VonkParameterException if MustHandle arguments are not handled.

```
try
{
    args.CheckHandled("Arguments must all be handled in MyService");
}
```

(continues on next page)

(continued from previous page)

```
catch (VonkParameterException vpe)
{
    //report it in the vonkContext.Response.Outcome
}
```

6. Inspect the number of the results to check whether anything was found. If so, you can enumerate the results or process the set as a whole, since `SearchResult` implements `IEnumerable<IResource>`.

```
if (searchResult.TotalCount > 0)
{
    foreach(var resource in searchResult)
    { ... }
}
```

## Implement ISearchRepository

Implementing `ISearchRepository` is only needed in a Facade.

The general pattern for implementing `ISearchRepository` is:

1. For each of the `IArguments` in the `IArgumentCollection`:
  1. If you support the argument, translate it to a ‘where’ clause on your repository. If your backend is another Web API, this could have the form of a querystring.
  2. Call `IArgument.Handled()` to update its status. There is also `.Warning()` and `.Error()` when something is wrong with the argument. If you simply don’t support the argument, you can leave the status to ‘Unhandled’.
  3. Pay special attention to the `_count` and `_skip` arguments for proper paging.
2. ‘AND’ all the arguments together, e.g. forming a database query or complete querystring.
3. Issue the query to your repository and await the results (await used intentionally: this should be done asynchronously).
4. For each of the resulting records or objects: map them to matching FHIR resources, either by:
  1. Creating POCO’s:

```
var result = new Patient() { /*fill in from the source object*/ };
return result.ToIResource(); //InformationModel implied by the assembly of
↳class Patient
```

2. or by crafting `SourceNodes`:

```
var result = SourceNode.Resource("Patient", SourceNode.Valued("id", /* id from
↳source */), SourceNode.Node("meta", SourceNode.Valued("versionId", "v1"), ....
↳), ...))
return result.ToIResource(VonkConstants.Model.FhirR3 /* or FhirR4 */);
```

5. Combine the mapped resources into a `SearchResult`:

```
return new SearchResult(resources, pagesize, totalCount, skip);
```

- `pagesize`: should be the value of the `_count` argument, unless you changed it for some reason.
- `totalCount`: total number of results, if there are more than you are returning right now.



- skip: number of results skipped in this set (if you are serving page x of y).

For a Facade on a relational database we provide a starting point with `Vonk.Facade.Relational.SearchRepository`. Follow the exercise in *Exercise: Build your first Facade* to see how that is done.

## IResourceChangeRepository

`IResourceChangeRepository` defines methods to change resources in the repository:

```
public interface IResourceChangeRepository
{
    Task<IResource> Create(IResource input);
    Task<IResource> Update(ResourceKey original, IResource update);
    Task<IResource> Delete(ResourceKey toDelete, string informationModel);
    string NewId(string resourceType);
    string NewVersion(string resourceType, string resourceId);
}
```

`ResourceKey` is a simple struct to identify a resource by Type, Id and optionally VersionId.

## Using IResourceChangeRepository

You should hardly ever need to use the `IResourceChangeRepository`. It is used by the *Create*, *Update*, *Delete* and the conditional variations thereof.

Should you need to use it, the methods are fairly straightforward.

### method

Create

### description

Provide an `IResource` having an id and versionId. These can be obtained by calling `NewId` and `NewVersion`. The return value will contain a possibly updated `IResource`, if the implementation changed or added elements. To fill in all the metadata there is a convenient extension method on `IResourceChangeRepository`:

```
var withMetaInfo = changeRepository.EnsureMeta(resource); //Will keep
↳existing id, version and lastUpdated and fill in if missing.
//EnsureMeta also exists as extension method on IResource - that uses
↳Guids for id and version.
var createdResource = await changeRepository.Create(withMetaInfo);
```

### method

Update

### description

Assert that a resource exists that can be updated. If not, use `Create`, otherwise go for `Update`.

```
var existingKey = new ResourceKey(resourceType, resourceId);
var args = existingKey.ToArguments(true);
var args = args.AddCount(0); //We don't need the actual result - just want
↳to know whether it is there.
var options = SearchOptions.Latest(vonkContext.ServerBase, VonkInteraction.
↳type_search, InformationModel: null); //search across informationmodels,
↳we expect ids to be unique.
```

(continues on next page)

(continued from previous page)

```

var exists = (await searchRepository.Search(args, options)).TotalCount = 1;
↪ //Take care of < 1 or > 1 matches

var withMetaInfo = changeRepository.EnsureMeta(resource, KeepExisting.Id);↪
↪ //Will keep existing id and provide fresh version and lastUpdated.
var updatedResource = await changeRepository.Update(existingKey,↪
↪ withMetaInfo);

```

**method**

Delete

**description**

Delete the resource that matches the provided key and informationModel. Returns the resource that was deleted.

```

var existingKey = new ResourceKey(resourceType, resourceId);
var deletedResource = await changeRepository(existingKey, vonkContext.
↪ InformationModel);

```

**method**

NewId

**description**

Get a new Id value generated by the repository (e.g. when the repository wants to use a sequence generator or ids in a specific format). Generally used through the extension method `IRResourceChangeRepository.EnsureMeta(IRResource resource, KeepExisting keepExisting)`, see `Create` above.

**method**

NewVersion

**description**

Get a new Version value generated by the repository (e.g. when the repository wants to use a sequence generator or ids in a specific format). The repository may want to base the version on the id, therefore the Id is passed as an argument. Generally used through the extension method `IRResourceChangeRepository.EnsureMeta(IRResource resource, KeepExisting keepExisting)`, see `Create` above.

## Implement `IRResourceChangeRepository`

Implementing `IRResourceChangeRepository` is only needed in a Facade that wants to provide write-access to the underlying repository.

For all three methods, you will have to map data from FHIR resources to your internal data structures and back.

Note that you also need to implement `ISearchRepository` to support the `Create` and `Update` plugins and of course the conditional variants of those.

### 11.4.7 Constructing bundles

In a Plugin or Facade you may need to construct a bundle from a set of resources, e.g. a SearchResult (see [here](#)). There are two ways of doing this: with the Bundle POCO or with SourceNodes.

#### Bundle POCO

This is fairly straightforward. Create a new Bundle object, and fill its properties, iterating over the set of resources that you have. The code looks like this:

```
using Hl7.Fhir.Model; //Either from Hl7.Fhir.Core.Stu3 or Hl7.Fhir.Core.R4

...

var searchResult = await searchRepository.Search(args, options)
var bundle = new Bundle() { Type = Bundle.BundleType.Searchset }; // Type is required

foreach (var resource in searchResult)
{
    bundle.Entry.Add
    {
        new BundleEntryComponent
        {
            Resource = resource.ToPoco<Resource>();
            // fill in any other details like Request or Response.
        }
    }
}

//fill in details of the bundle as a whole, like Meta or Identifier.
```

The limitation of this is that you are bound to either STU3 or R4, and that also implies that you cannot include *Custom Resources* in the bundle.

#### Bundle from SourceNodes

ISourceNode is from Hl7.Fhir.ElementModel and not tied to a specific FHIR version, and Firely Server can serialize ISourceNode provided that the right StructureDefinition is available in the Administration API - which is the case for Bundle by default.

You start by creating the Bundle itself using the class SourceNode that allows for construction of ISourceNode nodes.

```
using Hl7.Fhir.ElementModel;

...

var bundleNode = SourceNode.Resource("Bundle", "Bundle", SourceNode.Valued("type",
↪ "document"));

//choose type as one of the bundle types from the spec, see http://hl7.org/fhir/R4/
↪ bundle-definitions.html#Bundle.type
```

Then you can add elements to the bundle itself that are not in the entries of the bundle. Like an identifier:

```
var identifier = SourceNode.Node("identifier");
identifier.Add(SourceNode.Valued("system", "urn:ietf:rfc:3986"));
identifier.Add(SourceNode.Valued("value", Guid.NewGuid().ToString()));
bundleNode.Add(identifier);
```

Then you can turn this into the helper class `GenericBundle` that provides several helper methods on an `ISourceNode` that is known to be a `Bundle`.

```
var documentBundle = GenericBundle.FromBundle(bundleNode);
documentBundle = documentBundle.Meta(Guid.NewGuid().ToString(), DateTimeOffset.Now);
```

Maybe you already saw an alternative way of adding the identifier in the intellisense by now:

```
documentBundle = documentBundle.Identifier("urn:ietf:rfc:3986", Guid.NewGuid().
↳ ToString());
```

Note that you always have to continue with the *result* of the modifying function. All these functions act on `ISourceNode` and that is immutable, so you get a new instance with the changes applied as a return value.

Now you have the skeleton of the `Bundle`, it is ready to add entries with resources to it.

```
IResource resourceForDocument = ... ; //Get or construct a resource that is one of the
↳ entries of the Bundle.
documentBundle = documentBundle.AddEntry(resourceForDocument, resourceForDocument.Key().
↳ ToRelativeUri());
```

Other extensions methods available on `GenericBundle`:

```
public static GenericBundle Total(this GenericBundle bundle, int total)
public static GenericBundle AddLink(this GenericBundle bundle, string relation, string
↳ uri)
public static GenericBundle Links(this GenericBundle bundle, Dictionary<string, string>
↳ links)
```

### Search result bundles

Usually you don't need to construct a searchset bundle yourself, since the `SearchService` takes care of that when a search is issued on the FHIR endpoint. But should you want to do it in a custom operation, then the methods for doing so are at your disposal.

To help construct a bundle of type 'searchset', there is a special kind of bundle class `SearchBundle`. Create the sourcenode for the bundle as above. Then instead of creating a `GenericBundle`, turn it into a `SearchBundle`:

```
var searchBundle = bundleNode.ToSearchBundle();
```

Now you can use various methods to add entries for matches, includes or an `OperationOutcome`:

```
//SearchBundle methods
public SearchBundle AddMatch(ISourceNode resource, string fullUrl, string score = null)
public SearchBundle AddInclude(ISourceNode resource, string fullUrl, string score = null)
public SearchBundle AddOutcome(ISourceNode outcome, string fullUrl, string score = null)

//Extension methods
```

(continues on next page)

(continued from previous page)

```

public static SearchBundle ToSearchBundle(this IEnumerable<SearchInfo> searchInfos,
    ↪ string informationModel)
public static SearchBundle ToSearchBundle(this IEnumerable<ISourceNode> resources,
    ↪ string searchMode, string informationModel)
public static SearchBundle ToSearchBundle(this IEnumerable<ITypedElement> resources,
    ↪ string searchMode, string informationModel)

```

The SearchInfo struct essentially captures all the information that goes into an entry of a searchset bundle:

```

public struct SearchInfo
{
    public SearchInfo(ISourceNode resource, string mode = SearchMode.match, string
    ↪ fullUrl = null, string score = null)

    public string Mode { get; }
    public ISourceNode Resource { get; }
    public string FullUrl { get; }
    public string Score { get; }
}

```

Using all this to turn the SearchResult returned from the ISearchRepository.Search() method into a bundle looks like this (using the second extension method above):

```

var bundle = searchResult
    .ToSearchBundle(SearchMode.match, vonkContext.InformationModel)
    //informationModel is needed because bundle has slight differences between STU3
    ↪ and R4
    .Total(searchResult.Page.TotalCount)
    //Total is defined on GenericBundle
    .Links(searchResult.Page.PagingLinks(vonkContext));
    //Links is defined on GenericBundle
return bundle.ToIResource(vonkContext.InformationModel).EnsureMeta();

```

## 11.4.8 Access Control in Plugins and Facades

The *Access Control feature* is also available to users of a Firely Server Facade or Firely Server Plugins. You can use the default implementation based on SMART on FHIR, or provide an implementation of your own.

### Access control implementation

The access control engine is programmed using interfaces for which you can provide your own implementation. Because we think the model behind SMART on FHIR covers many cases, these interfaces are loosely modelled after it. The important interfaces and class are:

Table 1: Access Control interfaces

Interface / Class	Description
IAuthorization	Defines whether your are allowed to read or write a type of resource. Abstraction of the concept of a scope (like user/Observation.read) in SMART
ICompartment	Confines the user to a compartment, expressed as a combination of a CompartmentDefinition and a search argument. Abstraction of the concept of a launch context (like patient=123) in SMART
IReadAuthorizer	Calculates access control for a type of resource given an instance of IAuthorization and/or ICompartment
IWriteAuthorizer	Calculates access control for writing a new (version of a) resource given an instance of IAuthorization and/or ICompartment
AuthorizationResult	Return value of IReadAuthorizer and IWriteAuthorizer methods. It expresses whether you are authorized at all, and if so - under which conditions. These conditions are expressed as search arguments.

## **IReadAuthorizer**

Provides two methods to check authorization for reading types of resources.

- AuthorizeRead
- AuthorizeReadAnyType

The latter is only used if a system wide search is performed, without a `_type` parameter. In that case it is not efficient to call the first method for every supported resourcetype.

The input of these operations is an IAuthorization and an ICompartment. The result is an AuthorizationResult. With this class you can return:

- simply true or false
- extra search arguments to add to the search query in order to confine the search to those resources the user is allowed to read.

The AuthorizationResult Filters member is a collection of IArgumentCollections. Arguments within a collection will be AND'ed together. Multiple collections will be OR'ed together.

## **IWriteAuthorizer**

Provides one method to assess whether the user is allowed to write a resource. Input is again IAuthorization and ICompartment, but also IResource - the resource that is to be written - and an Uri called 'serverBase'. The 'serverBase' parameter is primarily provided because it is required to perform a search on the ISearchRepository interface. The IAuthorization instance can be used to decide whether the user is allowed to write resources of the given resourcetype at all. The ICompartment can be used to search in the database whether the to-be-written resource is linked to the current compartment.

## 11.5 Architecture

### 11.5.1 Pipeline and middleware components

Firely Server is built upon ASP.NET Core and its modular setup of a pipeline of [middleware](#). It also leverages the [dependency injection](#) pattern that is central to ASP.NET Core. If you are not familiar with these concepts, please review them first.

Firely Server's core plugins provide middleware components that handles the interactions of the [FHIR RESTful API](#). These components are placed in the pipeline. The pipeline starts with a few general components that interpret the incoming http request and determine which interaction is being called. This information is fed to the rest of the pipeline as the Request property of an `IVonkContext` object, analogous to the `HttpContext` and its Request property of ASP.NET Core itself. The `IVonkContext` also provides the Response property that is used to communicate the status, resource and/or `OperationOutcome` resulting from the interaction. On the way back through the pipeline the counterparts of the interpreting middleware translate the `IVonkContext` Response back to an http response, conforming to the FHIR specification. The rest of the pipeline consists of middleware components that each fulfill one of the interactions. So there is one for read, one for create, for search et cetera. These components are just regular ASP.Net Core Middleware components, except that they have access to the `IVonkContext`, and act on that.

#### Adding custom middleware

Using Firely Server Components you can add your own middleware anywhere in the pipeline. It can be standard ASP.NET Core middleware - having nothing to do with FHIR - or middleware acting on the `IVonkContext`, e.g. to implement a custom operation. Firely Server also provides convenience methods to register your middleware as one that handles a FHIR interaction, including attributes to declare for which interaction and which resource types your middleware should be invoked. This is explained in [Firely Server Plugins](#).

### 11.5.2 Repository interfaces

Many of the FHIR interactions require access to a repository of resources. E.g. `create` must be able to store a resource, whereas `search` must be able to query resources and retrieve the results. In Firely Server, the middleware components that implement these interactions access the repository through interfaces. There are two different interfaces for different parts of the [FHIR RESTful API](#).

<code>IChangeRepository</code>	<i>//for create, update and delete</i>
<code>ISearchRepository</code>	<i>//for all types of search, read and history</i>

In many scenarios, read-only access is sufficient, and you only need to implement the `ISearchRepository`. In that implementation you can choose which of the search parameters you want to support, and whether to expose versions and deleted resources.

These interfaces enable you to implement a Firely Server Facade. And they enable us to support database engines as diverse as MongoDB, SQL Server and in-memory.

## Search

The [FHIR RESTful Search](#) is the most complex part of the [FHIR RESTful API](#). Firely Server is capable of interpreting the search and translating it to small query-bits irrespective of the actual repository implementation. When implementing the `ISearchRepository` you get full control over which parameters you support and how you support them. On the method `ISearchRepository.Search()`, you just get the list of arguments that make up the search, as an `IArgumentCollection`. If you decide to act on these raw arguments directly, you can. But if you want Firely Server to interpret the search, you can use the `QueryBuilderContext.CreateQuery` method that will decompose the `IArgumentCollection`, interpret and validate every argument in it and then call into the `IRepoQueryFactory.Filter` or `IRepoQueryFactory.ResultShape` method for each of the arguments.

So the `Filter` method is called with interpreted search arguments. E.g. `identifier=abc` will be provided as (in pseudocode) `Filter("identifier", TokenValue{code = "abc", noSystem = true})`. If no search parameter is defined for an argument, you still get the change to handle it. E.g. `myspecialparameter=qyz` will be provided as `Filter("myspecialparameter", RawValue{ value = "qyz" })`. This allows for easy extensibility without defining your own `SearchParameter` resources, and is suitable for adding parameters that map to non-FHIR structures in your backend system. Note however that Firely Server also supports [Custom Search Parameters](#).

The `ResultShape` method is called when an argument is recognized as a ‘[Search result parameter](#)’, like `_sort` or `_include`.

### 11.5.3 Capabilities

A FHIR server has to express its capabilities in a `CapabilityStatement`, available on the `/metadata` endpoint. Firely Server’s capabilities are defined by the middleware components that make up its pipeline. Every component knows what interaction it adds to the capabilities. Therefore, we keep that information close to the component itself. Typically, every component has an implementation of `ICapabilityStatementContributor`, in which it gets access to the `ICapabilityStatementBuilder`. The latter provides methods to add information to the `CapabilityStatement` without having to worry about what is already added by other components or the order of execution.

## 11.6 Dependencies of Firely Server and their licenses

Firely Server is mainly built using libraries from Microsoft .Net Core and ASP.NET Core, along with a limited list of other libraries. This is the full list of direct dependencies that Firely Server has on other libraries, along with their licenses.

This list uses the NuGet package names (or prefixes of them) so you can easily lookup further details of those packages on [NuGet.org](#) if needed.

1. Microsoft.AspNetCore.\* - Apache 2.0
2. Microsoft.ApplicationInsights.\* - MIT
3. Microsoft.Bcl.AsyncInterfaces - MIT
4. Microsoft.EntityFrameworkCore.\* - MIT
5. Microsoft.Extensions.\* - MIT
6. Microsoft.AspNetCore.WebApi.Client - [MS-.NET-Library License](#)
7. System.Interactive.Async - MIT
8. Microsoft.Data.SqlClient - MIT
9. Microsoft.CSharp - MIT
10. System.Interactive.Async - MIT



11. System.Text.Json - MIT
12. System.\* - [MS-.NET-Library License](#)
13. NETStandard.Library - MIT
14. Newtonsoft.Json - MIT
15. IdentityModel.\* - Apache 2.0
16. IdentityServer4.AccessTokenValidation - Apache 2.0
17. IPNetwork2 - BSD 2-Clause “Simplified” License
18. Quartz - Apache 2.0
19. Serilog(\*) - Apache-2.0
20. LinkKit.Microsoft.EntityFrameworkCore - MIT
21. HL7.Fhir.\* - Firely OSS license (see below)
22. Fhir.Metrics - as HL7.Fhir
23. Simplifier.Licensing - as HL7.Fhir
24. Dapper - Apache 2.0
25. SqlKata.\* - MIT

MongoDB:

1. MongoDB.\* - Apache 2.0

For unittesting:

1. XUnit - Apache 2.0
2. Moq - BSD 3
3. FluentAssertions - Apache 2.0
4. Microsoft.NETCore.Platforms - MIT
5. Microsoft.NET.Test.Sdk - [MS-.NET-Library License](#)
6. System.Reactive - MIT
7. coverlet.collector - MIT
8. WireMock.Net - Apache 2.0

### 11.6.1 Firely OSS License

Firely Server relies on the reference .NET FHIR library: HL7.Fhir.\*, also created and maintained by Firely. The license is this (as stated in the [LICENSE](#) file:

Copyright (c) 2013-2020, HL7, Firely ([info@fire.ly](mailto:info@fire.ly)), Microsoft Open Technologies and contributors. See the file [CONTRIBUTORS](#) for details

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Firely nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This section contains documentation about tools that can improve your Firely Server experience.

## 12.1 Firely Server Ingest (FSI)

**Firely Server Ingest (FSI)** is a CLI application designed to optimize massive resource ingestion into a Firely Server instance. In contrast to resource ingestion by sending HTTP requests to a running Firely Server instance, this tool writes data directly to the underlying FS database which increases the throughput significantly.

The tool supports ingestion into SQL Server and MongoDB databases.

### 12.1.1 Installation

To install the tool, you first need to have .NET Core SDK v6.x installed on your computer. You can download it [here](#). Once it is installed, execute the following command:

```
dotnet tool install --global Firely.Server.Ingest
```

The command above will install FSI from this [NuGet package](#).

---

**Note:** Make sure that the dotnet tools directory is added to your path, this makes it possible to run the `fsi` command from any directory.

- Please find more information on where globally installed tools are located in [this article](#).
  - For Linux and Mac, make sure you add your `.profile` or `.bash_profile` to your path.
- 

### 12.1.2 General usage

**Attention:**

- Firely Server instances targeting the same database will be impacted severely by the workload that FSI puts on the database. We advise to stop the Firely Server instances while the import is performed.
- Only one instance of FSI per database should be run at a time. FSI can utilize all the cores on the machine it is run on, and insert data over several connections to the database in parallel. Multiple instances would probably cause congestion in the database.

### Prerequisites

The tool requires that the target database already exists and contains all required indexes and tables (for SQL Server). If you don't have a database with the schema yet, you first need to run the Firely Server at least once as described in the articles *Using SQL server* and *Using MongoDB*.

### Input files formats

FSI supports the following input file formats:

- FHIR *collection* bundles stored in \*.json files, and
- \*.ndjson files where each line contains a separate FHIR resource in JSON format.

### After the import

After ingesting massive amount of data, it is important to make sure the SQL Server indexes are in good shape. You can read more on this topic here: *Index Maintenance*.

### 12.1.3 Arguments

The execution of FSI can be configured using input parameters. These parameters can be supplied either as CLI arguments or specified in the file `appsettings.instance.json` which must be created in the same directory as the `fsi` executable.

If you want to specify input parameters in the file, you can use the snippet below as a base for your `appsettings.instance.json`. In this case, you need to update the values that you want to set yourself and delete all other records.

```
{
  "source": "./fsi-source", //valid directory
  "limit": -1,
  "fhirVersion": "R4",
  "license": "C:\\data\\deploy\\vonk\\license\\performance-test-license.json",
  "updateExistingResources": true,
  "databaseType": "SQL",
  "haltOnError": false,
  "convertAbsoluteUrlsToRelative": []

  "sqlserver": {
    "connectionString": "<connectionstring to the Firely Server SQL Server database>",
    "saveParallel": 2,
    "queryExistenceParallel": 4,
    "batchSize": 500,
    "commandTimeout": 60, //seconds
  },

  "mongodb": {
    "entryCollection": "vonkentries",
    "connectionString": "<connectionstring to the Firely Server MongoDB database>",
    "saveParallel": 2,
    "queryExistenceParallel": 4,
    "batchSize": 500
  },
}
```

(continues on next page)

(continued from previous page)

```

"workflow": { //-1 = unbounded
  "readParallel": 3,
  "readBufferSize": 200,
  "metaParallel": 1,
  "metaBufferSize": 50,
  "typeParallel": 4,
  "typeBufferSize": 50,
  "absoluteToRelativeParallel": 1,
  "absoluteToRelativeBufferSize": 50,
  "indexParallel": -1, //this is usually the most time consuming process - give it as
↳much CPU time as possible.
  "indexBufferSize": 50,
  "maxActiveResources": 15000
}
}

```

## Supported arguments

CLI argument	Appsettings parameter name	Required	Description
--settings <settingsJsonFile>			Custom
-f, --fhir-version <R3 R4>	fhirVersion		FHIR v
-s, --source <source>	source	yes	Input d
-l, --limit <limit>	limit		Limit t
--license <license>	license	yes	Firely v
--update-existing-resources <true false onlyIfNewer>	updateExistingResources		When
--dbType <MongoDb SQL>	databaseType		Specifi
--haltOnError <true false>	haltOnError		When
--convertAbsoluteUrlsToRelative <[array of values]>	convertAbsoluteUrlsToRelative		Conver
--mongoCollection <mongoCollection>	mongodb/entryCollection		Collect
--mongoConnectionString <connectionstring>	mongodb/connectionString	yes	Conne
--mongoPar <mongoPar>	mongodb/saveParallel		The nu
--mongoExistQryPar <mongoExistQryPar>	mongodb/queryExistenceParallel		The nu
--mongoBatch <mongoBatch>	mongodb/batchSize		The nu
-c, --connectionstring <connectionstring>	sqlServer/connectionString	yes	Conne
--sqlPar <sqlPar>	sqlServer/saveParallel		The nu
--sqlBatch <sqlBatch>	sqlServer/saveBatchSize		The nu
--sqlTimeout <sqlTimeout>	sqlServer/commandTimeOut		The tin
--sqlExistQryPar <sqlExistQryPar>	sqlserver/queryExistenceParallel		The nu
--readPar <readPar>	workflow/readParallel		Numbe
--readBuffer <readBuffer>	workflow/readBufferSize		Numbe
--metaPar <metaPar>	workflow/metaParallel		Numbe
--metaBuffer <metaBuffer>	workflow/metaBufferSize		Numbe
--typePar <typePar>	workflow/typeParallel		Numbe
--typeBuffer <typeBuffer>	workflow/typeBufferSize		Numbe
--absRelPar <absRelPar>	workflow/ absoluteToRelativeParallel		Numbe
--absRelBuffer <absRelBuffer>	workflow/ absoluteToRelativeBufferSize		Numbe
--indexPar <indexPar>	workflow/indexParallel		Numbe
--indexBuffer <indexBuffer>	workflow/indexBufferSize		Numbe

<code>--maxActiveRes &lt;maxActiveRes&gt;</code>	workflow/maxActiveResources		Maxim
<code>--version</code>			Show v
<code>-, -h, --help</code>			Show h

### 12.1.4 Examples

Specify a custom settings file `/path/to/your/custom/settings/appsettings.instance.json`.

```
fsi --settings ./path/to/your/custom/settings/appsettings.instance.json
```

**Note:** If `--settings` is omitted, FSI searches following folders sequentially and tries to find `appsettings.instance.json`. The first occurrence will be used if FSI finds one, otherwise the default `appsettings.json` will be used.

- Current launched folder e.g. `C:\Users\Bob\Desktop`
- FSI installation folder e.g. `C:\Users\Bob\.dotnet\tools`
- FSI installation dll folder e.g. `C:\Users\Bob\.dotnet\tools\.store\firely.server.ingest\version\firely.server.ingest\version\tools\net6.0\any`

Run the import for files located in directory `/path/to/your/input/files` and its subdirectories using license file `/path/to/your/license/fsi-license.json` targeting the database defined by the connection string. In case a resource being imported already exists in the target database, it gets skipped.

```
fsi \
-s ./path/to/your/input/files \
--license /path/to/your/license/fsi-license.json \
-c 'Initial Catalog=VonkData;Data Source=server.hostname,1433;User ID=username;
↳Password=PaSSSword!' \
--update-existing-resources false
```

Same as above but if a resource being imported already exists in the target database, it gets updated. The old resource gets preserved as a historical record.

```
fsi \
-s ./path/to/your/input/files \
--license /path/to/your/license/fsi-license.json \
-c 'Initial Catalog=VonkData;Data Source=server.hostname,1433;User ID=username;
↳Password=PaSSSword!'
```

Same as above but targeting a MongoDB database.

```
fsi \
--dbType MongoDB
-s ./path/to/your/input/files \
--license /path/to/your/license/fsi-license.json \
--mongoConnectionString 'mongodb://username:password@localhost:27017/vonkdata'
```

## 12.1.5 Monitoring

### Logs

When importing the data, it is handy to have the logging enabled, as it would capture any issues if they occur. By default, the log messages are written both to the console window and to the log files in the %temp% directory.

You can configure the log settings the same way as you do for Firely Server: [Log settings](#).

### Performance counters

You can get insights into the tool performance by means of performance counters. There are many ways to monitor the performance counters. One of the options is using [dotnet-counters](#).

To monitor the counters for FSI, you can execute the following command:

```
dotnet-counters monitor --counters 'System.Runtime','FSI Processing' --process-id
<process_id>
```

where `<process_id>` is the PID of the running FSI tool.

---

**Note:** If you think the ingestion process is going too slow for your amount of data and the hardware specifications, please [contact us](#) for advice.

---

## 12.1.6 Known issues

- FSI does not support scenarios where resources of different FHIR versions are stored in the same database;
- When importing data from large \*.ndjson files, the memory consumption may be quite high.
- When importing STU3 resources, the field `Patient.deceased` will always be set to `true` if it exists. This is caused by an error in the FHIR STU3 specification. In case you would like to use FSI with STU3 resources, please [contact us](#).

## 12.1.7 Licensing

The application is licensed separately from the core Firely Server distribution. Please [contact](#) Firely to get the license.

Your license already permits the usage of FSI if it contains `http://fire.ly/vonk/plugins/bulk-data-import`.

## 12.1.8 Release notes

### Release 1.4.0, October 6th, 2022

- Added new setting `convertAbsoluteUrlsToRelative` which is an array of server URL base values. This feature converts absolute URL references to relative references for the given server URL base array. Example: Setting of `http://example.org/R4` will convert an absolute URL `http://example.org/R4/Patient/123` to relative as `Patient/123`.

- Added a new mode `onlyIfNewer` for option `--update-existing-resources` (see the CLI options above)

---

**Note:** This option is currently supported only for SQL Server

---

- The setting `--useUcum` has been removed. From now on, all quantitative values get automatically canonicalized to UCUM values
- Indexing has been fixed for search parameters of type *reference* that index resource elements of type *uri*. The following SearchParameters were affected by the bug:
  - FHIR4: `ConceptMap-source-uri`, `ConceptMap-target-uri`, `PlanDefinition-definition`
  - STU3: `ImplementationGuide-resource`, `Provenance-agent`

Consider *re-indexing* your database for these search parameters if you use them.

---

**Note:** Please note that due to a mistake in the official STU3 specification, search parameters *ConceptMap-source-uri*, *ConceptMap-target-uri* still do not work as expected. The correct search parameter expressions would be *ConceptMap.source.as(uri)* and *ConceptMap.target.as(uri)* while the specification contains *ConceptMap.source.as(Uri)* and *ConceptMap.target.as(Uri)* respectively. The issue has been addressed in R4.

---

### Release 1.3.1

- Corrected an exception when multiple batch threads are processing and saving in parallel to SQL Server.

### Release 1.3.0

- Add configuration `haltOnError`. When `true`, the FSI will be stopped on a single error. Otherwise, it will log error and continue.
- Changed the serialization format of decimal from string to use the native decimal type in MongoDB to improve performance.
- Bugfix: Fixed `Money.currency` indexing for FHIR STU3 and R4

### Release 1.2.0

- Ability to provide a path to a custom `appsettings.json` file via a command-line argument (see *examples* above)
- Bugfix: ensure FSI uses all available values from the SQL PK-generating sequences when inserting data to the `vonk.entry` and `component` tables



## Release 1.1.0

- Feature: added support for MongoDB!
- Feature: added support for performance counters using dotnet-counters. See [Performance counters](#) on how to setup and use dotnet-counters.
- FSI has been upgraded to .NET 6. To install the tool, you first need to have .NET Core SDK v6.x installed on your computer. See [Installation](#) for more information.
- The Firely .NET SDK that FSI uses has been upgraded to 3.7.0. The release notes for the SDK v3.7.0 can be found [here](#).
- Multiple smaller fixes to improve reliability and performance of the tool.

## Release 1.0.0

- First public release
- Performance: optimized memory consumption (especially, when reading large \*.ndjson files)
- Feature: quantitative values can be automatically canonicalized to UCUM values (see `-useUcum` CLI option)
- Multiple smaller fixes to improve reliability and performance of the tool

# 12.2 Firely Auth

This section contains documentation about Firely Auth.

## 12.2.1 Firely Auth release notes

### Release 3.0.0, December 2022

This is the first public release of Firely Auth, providing support for SMART on FHIR v1 and v2 and a SQL Server user store.

## 12.2.2 Firely Auth Introduction

**Firely Auth** is an authentication and authorization server that implements the [SMART on FHIR](#) authentication flows. It is an add-on to Firely Server.

Several implementation guides and scenarios require the use of SMART on FHIR:

- Bulk Data Export
- (g)(10) certification
- Inferno testing

Firely Auth can be used to fulfill all of them.

There are various ways to get to know Firely Auth and the way it works together with Firely Server. As an introduction we'll setup both to make an example Postman collection work. In the configuration section, we discuss the configuration that is possible both in Firely Auth itself and in Firely Server.

### Getting started

#### Step 1 - Software

Firely Auth is distributed as .NET Core 6 binaries and in a Docker image. For this introduction we will use the binaries.

1. Install .NET Core 6 Runtime

#. You can request a copy of the binaries from Firely through our [contact page](#) .. #. Download the zip file with Firely Auth binaries from [the download server](#) #. Extract the zip to a location from where you are allowed to execute the program. We will call this the 'bin-directory'

#### Step 2 - License

Firely Auth is licensed, like all plugins and add-ons of Firely Server. It uses the same license file as the Firely Server instance it works with. Firely Auth requires this token to be present in the license file: <http://fire.ly/server/auth>. If you don't have this in your license file yet, you probably need to acquire Firely Auth first. Please [Contact us](#) for that.

By default Firely Auth will look for a license file named `firely-server-license.json`, adjacent to the `Firely.IdentityServer.Core.exe`. You can adjust the location of the license file in the configuration settings, see [License](#). Additionally you will have to place a file called `Duende_License.key` also adjacent to the `Firely.IdentityServer.Core.exe`. Please note that the path to this file cannot be configured.

With the license in place, you can start Firely Auth by running:

```
> ./Firely.IdentityServer.Core.exe
```

And you can access it with a browser on <https://localhost:5001>. It will use a self-signed certificate by default, for which your browser will warn you. Accept the risk and proceed to the website.

Firely Auth will present you with a login screen. But what user to login as?

#### Step 3 - Users

You need to add at least one user to Firely Auth. Firely Auth supports two types of stores for users: In memory and SQL Server. For this introduction we will configure a user in the In memory store.

1. Go to the bin-directory
2. Open `appsettings.default.json`
3. Look for the section `UserStore`
4. By default the `Type` is already set to `InMemory`.
5. Now add to the `InMemory:AllowedUsers` array:

```
{  
  "Username": "alice",  
  "Password": "AlicePassword1!",  
  "AdditionalClaims": []  
}
```

We won't add any claims yet - see [User store](#) to read up on how they work.

If you did this, you can stop Firely Auth again in the command window you started it in (Ctrl+C), and start it again. You should now be able to login as `alice` with the password as configured above.

This time, Firely Auth will tell you that there are no clients configured that can access data on your behalf.

## Step 4 - Clients

The concept of OAuth2 in general and SMART on FHIR in particular is that a client (an app, a website) can access data on your behalf. This means that Firely Auth must know these clients upfront. For each client several values need to be configured. For this introduction we will add Postman as a client, so you can test requests without actually building a client yourself. We'll just provide the correct settings here. The settings are documented in detail on [Clients](#)

**Note:** Making Postman trust the self-signed certificate of Firely Auth is outside the scope of this tutorial. For the purpose of this tutorial you can instruct Postman to not check SSL certificates.

```
"ClientRegistrationConfig": {
  "AllowedClients": [
    {
      "ClientId": "Jv3nZkaxN36ucP33",
      "ClientName": "Postman",
      "Description": "Postman API testing tool",
      "Enabled": true,
      "RequireConsent": true,
      "RedirectUri": ["https://www.getpostman.com/oauth2/callback", "https://
↪oauth.pstmn.io/v1/callback"],
      "ClientSecrets": [{"SecretType": "SharedSecret", "Secret": "re4&ih)+HQu~w"}],
      "AllowedGrantTypes": ["client_credentials"],
      "AllowedSmartLegacyActions": [],
      "AllowedSmartActions": ["c", "r", "u", "d", "s"],
      "AllowedSmartSubjects": [ "patient", "user", "system"],
      "AlwaysIncludeUserClaimsInIdToken": true,
      "RequirePkce": false,
      "AllowOfflineAccess": false,
      "AllowOnlineAccess": false,
      "AllowFirelySpecialScopes": true,
      "RequireClientSecret": true,
      "LaunchIds": []
    }
  ]
}
```

The values for ClientId and ClientSecrets.Secret are randomly generated. You are recommended to generate your own values.

We will use Postman to issue a request for an Access Token. For this we created a collection 'Firely Auth docs', and we will set the Authorization for the collection as a whole. That way the authorization can be reused for all requests in the collection. Click 'Get New Access Token' and you'll be taken to the login page of Firely Auth. If you are still logged in since step 3, you will be authorized immediately.

If the authorization request fails, check both the Postman console and the Firely Auth logging for a clue.

In the Authorization tab of the collection, set up the values according to the client settings above, see the image below. Note that we also set the Audience in the Advanced Settings to the default value **Firely Server**. This corresponds to settings discussed below.

## Firely Auth docs

Share For

Authorization • Pre-request Script Tests Variables Runs

This authorization method will be used for every request in this collection. You can override this by specifying one in the request.

Type OAuth 2.0

The authorization data will be automatically generated when you send the request. Learn more about [authorization](#)

Add auth data to Request Headers

Heads up! These parameters hold sensitive data. To keep this data secure while working in a collaborative environment, we recommend using variables. Learn more about [variables](#)

### Current Token

Access Token Available Tokens

Access Token

Header Prefix Bearer

### Configure New Token

Configuration Options • Advanced Options

Token Name Pm\_FA\_token1

Grant Type Authorization Code

Callback URL <https://oauth.pstmn.io/v1/callback>

☒ Authorize using browser

Auth URL [{{ids}}/connect/authorize](#)

Access Token URL [{{ids}}/connect/token](#)

Client ID Jv3nZkaxN36ucP33

Client Secret re4%26ih)%2BHQ~w

Scope openid fhirUser user/\*.rs

State State

Client Authentication Send as Basic Auth header

Clear cookies

Get New Access Token

## Configure New Token

Configuration Options ●

Advanced Options ●

Resource ⓘ

Resource

Add another resource

Audience ⓘ

Firely Server

Add another audience

Get New Access Token

**Note:** Encoding the secret The client secret as set in the `ClientRegistrationConfig` contains characters that must be URI-encoded. For secure secrets this may happen. In Postman, select the client secret string, right-click and choose “EncodeURIComponent”. For other clients you may use any other URI encoding tool, or encode it in your code before sending the access token request.

Client ID ⓘ

Jv3nZkaxN36ucP33

Client Secret ⓘ

re4%26ih)%2BHQ~w

Scope ⓘ

user/Patient

Client Authentication

Send as Basic

Clear cookies ⓘ

Get New Access Token

Set: Firely Auth docs ▶

Set: Globals ▶

Undo

Redo

Cut

Copy

Paste

Select All

Find: re4%26ih)%2BHQ~w

EncodeURIComponent

DecodeURIComponent

### Step 5 - Connect Firely Server to Firely Auth

Ultimately the access token that we just retrieved is meant to get access to resources in Firely Server. To demonstrate that we will:

1. Set up Firely Server locally.
2. Adjust the settings to connect it to Firely Auth

Setting up Firely Server is described in [Getting Started](#). Please follow that instruction if you have not already done so. For this introduction you can use the default settings and repositories for both data and administration, being SQLite. We will adjust only 1 setting to more easily work with FHIR R4:

```
"InformationModel": {  
  "Default": "Fhir4.0", // information model to use when none is specified in either_  
↪mapping, the _format parameter or the ACCEPT header  
  "IncludeFhirVersion": ["Fhir4.0", "Fhir5.0"],  
  "Mapping": {  
    "Mode": "Path", // yourserver.org/r3 => FHIR STU3; yourserver.org/r4 => FHIR R4  
    "Map": {  
      "/R3": "Fhir3.0",  
      "/R4": "Fhir4.0"  
    }  
  }  
},
```

With this, we can use <base>/R4 to use FHIR R4 (see for background [Multiple versions of FHIR](#)).

Check that it runs without authorization before proceeding with the next step, by requesting the CapabilityStatement:

Firely Auth docs / Check FS w/o auth

GET {{BASE\_URL}}/R4/metadata Send

Params Authorization Headers (7) Body Pre-request Script Tests Settings Cookies

Type No Auth

This request does not use any authorization. Learn more about [authorization](#)

Body Cookies Headers (6) Test Results 200 OK 105 ms 1.94 MB Save Response

Pretty Raw Preview Visualize JSON

```

1 {
2   "resourceType": "CapabilityStatement",
3   "id": "46762bdd-bbfi-41ca-96f3-6189c6a78d2d",
4   "meta": {
5     "versionId": "048f47f1-1dde-4b26-aefe-9d389ce8e47a",
6     "lastUpdated": "2022-11-08T13:53:52.0091397+00:00"
7   },
8   "language": "en-US",
9   "url": "http://server.fire.ly/fhir/CapabilityStatement/FirelyServer",
10  "version": "1.0",
11  "name": "Firely Server 4.10.0 CapabilityStatement",
12  "status": "active",
13  "experimental": true,
14  "date": "2022-11-08T13:53:52.009618+00:00",
15  "publisher": "Firely",
16  "contact": [

```

To be able to test the next steps, add a few example resources by issuing a batch request (POST <base>/R4/) with this bundle (while authorization is still off). It contains two Patient resources and an Observation related to each of them.

Now we will connect Firely Server and Firely Auth. This requires mutual settings.

In **Firely Auth**:

```

"FhirServerConfig": {
  "Name": "Firely Server",
  "FHIR_BASE_URL": "http://localhost:4080"
},

```

The Name in this section serves two purposes:

- it acts as the username for accessing the token introspection point.
- it is used for the *aud* (Audience) claim in the access token supplied to the requesting app.

The FHIR\_BASE\_URL is the url on which Firely Server can be reached by the requesting app. It is used to turn the *fhirUser* claim (e.g. Patient/123) into a full url.

In **Firely Server**, all the settings are in the section *SmartAuthorizationOptions*

```

"SmartAuthorizationOptions": {
  "Enabled": true,

```

(continues on next page)

(continued from previous page)

```

"Filters": [
  {
    "FilterType": "Patient", //Filter on a Patient compartment if a 'patient' launch
    ↳scope is in the auth token
    "FilterArgument": "_id=#patient#" //... for the Patient that has an id matching
    ↳the value of that 'patient' launch scope
  }
],
"Authority": "https://localhost:5001",
"Audience": "Firely Server", //Has to match the value the Authority provides in the
↳audience claim.
"RequireHttpsToProvider": true, //You want this set to true (the default) in a
↳production environment!
"Protected": {
  "InstanceLevelInteractions": "read, vread, update, patch, delete, history,
↳conditional_delete, conditional_update, $validate, $meta, $meta-add, $meta-delete,
↳$export, $everything, $erase",
  "TypeLevelInteractions": "create, search, history, conditional_create, compartment_
↳type_search, $export, $lastn, $docref",
  "WholeSystemInteractions": "batch, transaction, history, search, compartment_system_
↳search, $export, $exportstatus, $exportfilerequest"
},
// "TokenIntrospection": {
//   "ClientId": "vonk",
//   "ClientSecret": "secret"
// },
"ShowAuthorizationPII": false,
//"AccessTokenScopeReplace": "-",
"SmartCapabilities": [
  "LaunchStandalone",
  "LaunchEhr",
  //"AuthorizePost",
  "ClientPublic",
  "ClientConfidentialSymmetric",
  //"ClientConfidentialAsymmetric",
  "SsoOpenidConnect",
  "ContextStandalonePatient",
  "ContextStandaloneEncounter",
  "ContextEhrPatient",
  "ContextEhrEncounter",
  "PermissionPatient",
  "PermissionUser",
  "PermissionOffline",
  "PermissionOnline",
  "PermissionV1",
  //"PermissionV2",
  "ContextStyle",
  "ContextBanner"
]
},

```

All settings are discussed in detail in *Firely Server*, and we'll focus on the connection with Firely Auth here:



- Authority: the address where Firely Auth can be reached.
- Audience: By default Firely Server, should match the FhirServerConfig.Name setting in Firely Auth and the requested aud in Postman.

Now we should be able to issue an authorized request to Firely Server with the token we requested on the collection in Step 4.

Firely Auth docs / Read Patient with auth

GET `{{BASE_URL}}/R4/Patient` **Send**

Params Auth Headers (10) Body Pre-req. Tests Settings Cookies

**Auth**

**type**

Inherit auth from parent

The authorization header will be automatically generated when you send the request. Learn more about [authorization](#).

This request is using OAuth 2.0 from collection [Firely Auth docs](#).

Body Cookies Headers (7) Test Results 200 OK 3.05 s 2.22 KB **Save Response**

Pretty Raw Preview Visualize JSON

```

1
2   "resourceType": "Bundle",
3   "type": "searchset",
4   "timestamp": "2022-11-14T10:07:56.975+01:00",
5   "meta": {
6     "lastUpdated": "2022-11-14T10:07:56.975+01:00",
7     "versionId": "e91d81b1-57a1-4b2e-a74a-21f154b8bdba"
8   },
9   "entry": [
10    {
11      "fullUrl": "http://localhost:4080/R4/Patient/example_eddie",
12      "search": {
13        "mode": "match"
14      },
15      "resource": { ...
38    },
39    {
40      "fullUrl": "http://localhost:4080/R4/Patient/example_fred",
41      "search": {
42        "mode": "match"
43      },
44      "resource": { ...
45      }
46    }
47  ],
48  "total": 2,
49  "link": [
50    {
51      "relation": "self",
52      "url": "http://localhost:4080/R4/Patient?_total=accurate&_count=10&_skip=0"
53    }
54  ],
55  "id": "f6697306-5c25-42b1-82e2-fe60c6b4fd3c"
56

```

### 12.2.3 Firely Auth Settings

Firely Auth can be configured extensively. This page lists all the settings and refers to some detail pages for some of the sections.

#### Settings files and variables

Just like Firely Server itself, Firely Auth features a hierarchy of settings files and variables. From lowest to highest priority:

- `appsettings.default.json` - This comes with the binaries (and in the Docker container) and contains sensible defaults for most settings. You can change this file, but it might accidentally be overwritten upon a new release. Instead, put your settings in one of the following places.
- `appsettings.instance.json` - This file is meant to override settings for this instance of Firely Auth. You can create this file yourself by copying (parts of) `appsettings.default.json`.
- Environment variables with the prefix `FIRELY_AUTH_` - Use environment variables to more easily configure settings from CI/CD pipelines, secure vaults etc.

Unsure how to name your variables? This works the same as with *Firely Server settings with Environment Variables*.

#### Sections

##### License

Use the License settings to set the location to the license file. A relative path is evaluated relative to the executable `Firely.IdentityServer.Core.exe`. You can use the same license file that came with Firely Server.

```
"License": {  
  "LicenseFile": "../firelyserver-license.json"  
},
```

##### Kestrel

Kestrel is the web server that is integrated into Firely Auth, just like for Firely Server. We recommend to run both behind a *reverse proxy*. Nevertheless you can control the settings for Kestrel.

```
"Kestrel": {  
  "Endpoints": {  
    "Http": {  
      "Url": "http://localhost:5100"  
    },  
    "HttpsFromPem": {  
      "Url": "https://localhost:5101",  
      "SslProtocols": [ "Tls12", "Tls13" ],  
      "Certificate": {  
        "Path": "cert.pem",  
        "KeyPath": "cert-key.pem"  
      }  
    }  
  }  
},
```

These settings are not Firely Auth specific, and you can read more about them in the [Microsoft documentation](#). In that documentation you can also read how to use the `Https` setting instead of `HttpsFromPem` to use a `.pfx` file for your SSL certificate.

## Firely Server

Firely Auth hands out SMART on FHIR access tokens to access resources on Firely Server. To make Firely Server known to Firely Auth, fill in the `FhirServerConfig`:

```
"FhirServerConfig": {
  "Name": "Firely Server",
  "FHIR_BASE_URL": "http://localhost:4080"
},
```

- **Name:** This name serves two purposes:
  - It can be added to the token as the value of the `aud` (audience) claim, if the client requests so. To have it accepted by Firely Server, set its `SmartAuthorizationOptions:Audience` setting to the same value.
  - It correlates with the clients allowed to access the token introspection endpoint (see below). Therefore it should match the value of `TokenIntrospectionConfig:TokenIntrospectionResources:Name`
- **FHIR\_BASE\_URL:** This also has two uses:
  - A token can have a claim in the form of `patient=<base>/Patient/123`, to define the compartment the client is restricted to. This url is used as the base part in that url, and should match the base url of Firely Server, as it is accessed by the client.
  - If an `aud` parameter is provided *in the authorize request*, it has to match this url. E.g. in Postman you can provide this parameter by adding it to the Auth URL, like this: `{{ids}}/connect/authorize?aud=http://localhost:4080` See the `aud` parameter in [SMART on FHIR authorization request](#)

## Token types

Define for each client what type of token it can request.

```
"TokenConfig": {
  "AccessTokenType": {
    "<ClientId>": "Jwt"
  }
},
```

- `<ClientId>` should match one of the clients defined in `ClientRegistrationConfig`.
- The value can be one of `Jwt` or `Reference`. `Jwt` means that this client will get self-contained Json Web Tokens. `Reference` means that this client will get reference tokens, that refer to the actual token kept in memory by Firely Auth. For more background see [reference token](#).

E.g. `"MySmartApp": "Reference"`

## Key management

```
"KeyManagementConfig": {
  "RSA_Config": {
    ///"RSA_JWK": "<JSON Web Key>", // JSON Web Key of type RSA
    "SupportedAlgorithms": [
      "RS256",
      "RS384",
      "RS512"
    ]
  },
  "EC_Config": {
    ///"JWK_ES256": "<JSON Web Key>", // JSON Web Key of type EC with crv P-256
    ///"JWK_ES384": "<JSON Web Key>", // JSON Web Key of type EC with crv P-384
    ///"JWK_ES512": "<JSON Web Key>", // JSON Web Key of type EC with crv P-512
    "SupportedAlgorithms": [
      "ES256",
      "ES384",
      "ES512"
    ]
  }
}
```

Firely Auth can work with multiple signature keys, used to sign access and other tokens.

- **RSA\_Config**: defines the RSA algorithms that are supported. In the config above all available algorithms are listed. Inferno tests require at least RS256 for all Single Patient tests, and for Bulk Data Export a RS384 or higher is needed.
  - **RSA\_JWK**: allows to provide a pre-generated JSON Web Key. If this is not provided, Firely Auth will generate a key.
  - **SupportedAlgorithms**: limit this list to the algorithms that you need in your setup. In the config above all available algorithms are listed.
- **EC\_Config**: defines the EC (Elliptic Curve) algorithms that are supported. Inferno tests for Bulk Data Export require support for EC keys.
  - **JWK\_ES\***: allows to provide a pre-generated JSON Web Key. If this is not provided, Firely Auth will generate a key for each of the supported algorithms.
  - **SupportedAlgorithms**: limit this list to the algorithms that you need in your setup. In the config above all available algorithms are listed.

Note that a single RSA key can be used for all supported algorithms. However, an EC key is tied to a specific algorithm, therefore you can supply a key for each of the algorithms.

For more background on JSON Web Keys see [RFC 7517](#).

## Token introspection

When using a *reference token*, Firely Server must verify the token with Firely Auth. Not just any system can ask for inspection though, therefore we list the systems that can with a name and a secret.

```
"TokenIntrospectionConfig": {
  "TokenIntrospectionResources": [{
    "Name": "Firely Server",
    "Secret": "<generate some hard to hack secret>"
  }]
},
```

This configuration is only needed if at least one *client* is configured to use reference tokens, see *Token types*.

## User store

A *user* must be able to authenticate to Firely Auth before granting permissions to a *client*. Therefore we register the users with Firely Auth. Firely Auth supports two types of stores: In memory and SQL Server.

For the InMemory store, the users and their passwords are listed in plain text in this configuration. This is useful for testing, but not recommended for production use.

The SqlServer store stores the users and their encrypted passwords in a MS SQL Server database. Also the *fhirUser* and *patient* claims for each user can be stored. See *SQL Server user store* for details on setting up the database.

```
"UserStore": {
  "Type": "InMemory", // InMemory | SqlServer
  "InMemory": {
    "AllowedUsers": [
      {
        "Username": "bob",
        "Password": "password",
        "AdditionalClaims": [
          {
            "Name": "patient",
            "Value": "Patient/a123"
          }
        ]
      }
    ]
  },
  "SqlServer": {
    "ConnectionString": "<connection string here>"
  }
},
```

- Type: select the type of store to use
- InMemory: settings for the InMemory store
  - AllowedUsers: list of users
  - Username: login for a user
  - Password: password for the user, in clear text

- **AdditionalClaims**: currently to be used for a single claim, to link the user to a Patient resource (and thereby to a Patient compartment), or a ‘user’ resource like a Practitioner in Firely Server.
  - \* **Name**: name of the claim, currently only `patient` and `fhirUser` are supported
  - \* **Value**: logical id of the related Patient or Practitioner resource (`Patient/id`) In the token this value will be expanded to an absolute url by prepending it with `FhirServerConfig.FHIR_BASE_URL` (see *Firely Server*).
- **SqlServer**: settings for the SQL Server store
  - **ConnectionString**: connection string to the SQL Server database where the users are to be stored. This database and the schema therein must be created beforehand with a script.

## Clients

The `ClientRegistrationConfig` is used to register the *clients* that are allowed to request access tokens from Firely Auth.

```
"ClientRegistrationConfig": {
  "AllowedClients": [
    {
      "ClientId": "Jv3nZkaxN36ucP33",
      "ClientName": "Postman",
      "Description": "Postman API testing tool",
      "Enabled": true,
      "RequireConsent": true,
      "RedirectUris": ["https://www.getpostman.com/oauth2/callback", "https://
↪oauth.pstmn.io/v1/callback"],
      "ClientSecrets": [{"SecretType": "SharedSecret", "Secret": "re4&ih)+HQu~w"}],
      ↪ // SharedSecret, JWK
      "AllowedGrantTypes": ["client_credentials", "authorization_code"],
      "AllowedSmartLegacyActions": [],
      "AllowedSmartActions": ["c", "r", "u", "d", "s"],
      "AllowedSmartSubjects": [ "patient", "user", "system"],
      "AlwaysIncludeUserClaimsInIdToken": true,
      "RequirePkce": false,
      "AllowOfflineAccess": false,
      "AllowOnlineAccess": false,
      "AllowFirelySpecialScopes": true,
      "RequireClientSecret": true,
      "LaunchIds": [],
      "RefreshTokenLifetime": "30"
    }
  ]
}
```

You register a *client* in the `AllowedClients` array. For each client you can configure these settings:

- **ClientId**: string: unique identifier for this client. It should be known to the client as well
- **ClientName**: string: human readable name for the client, it is shown on the consent page
- **Description**: string: human readable description of the client
- **Enabled**: true / false: simple switch to enable or disable a client (instead of removing it from the list)

- **RequireConsent:** true / false: when true, Firely Auth will show the user a page for consent to granting the requested scopes to the client, otherwise all requested and valid scopes will be granted automatically.
- **RedirectUris:** array of strings: url(s) on which Firely Auth will send the authorization code and access token. The client can specify one of the preregistered urls for a specific request.
- **ClientSecrets:** secrets can be of type `SharedSecret` or `JWK`. You can have multiple of each, so you can accept two secrets for a short period of time to support key rotation and an update window for the client. The `ClientSecrets` section is ignored if `RequireClientSecret` is set to `false`.
  - `SharedSecret`: `{"SecretType": "SharedSecret", "Secret": "<a secret string shared with the client>"}` - this can be used for either *client credentials* or *authorization code flow*, but only with a *confidential client*.
  - `JWK`: `{"SecretType": "JWK", "SecretUrl": "<JWKS url>"}` - where the JWKS url hosts a JSON Web Key Set that can be retrieved by Firely Auth, see also *JWK*.
  - `JWK`: `{"SecretType": "JWK", "Secret": "<JWK>"}` - where JWK is the contents of a *JWK*. Use this if the client cannot host a url with a JWKS. Use one entry for each key in the keyset. Note that the JWK json structure is embedded in a string, so you need to escape the quotes within the JWK. The url option above is recommended.
- **AllowedGrantTypes:** array of either or both `"client_credentials"` and `"authorization_code"`, referring to *client credentials* and *authorization code flow*. Use `client_credentials` only for a *confidential client*.
- **AllowedSmartLegacyActions:** Firely Auth can also still support SMART on FHIR v1, where the actions are read and write.
- **AllowedSmartActions:** Actions on resources that can be granted in SMART on FHIR v2: `c`, `r`, `u`, `d` and/or `s`, see *SMART on FHIR V2 scopes*
- **AllowedSmartSubjects:** Categories of 'subjects' to which resource actions can be granted. Can be `system`, `user` and/or `patient`
- **AlwaysIncludeUserClaimsInIdToken:** true / false: When requesting both an id token and access token, should the user claims always be added to the id token instead of requiring the client to use the `userinfo` endpoint. Default is false
- **Require PKCE:** true / false - see *PKCE*. true is recommended for a *public client* and can offer an extra layer of security for *confidential client*.
- **AllowOfflineAccess:** true / false - Whether app can request refresh tokens while the user is online, see *SMART on FHIR refresh tokens*
- **AllowOnlineAccess:** true / false - Whether app can request refresh tokens while the user is offline, see *SMART on FHIR refresh tokens*. A user is offline if he is logged out of Firely Auth, either manually or by expiration
- **AllowFirelySpecialScopes:** true / false - Allow app to request scopes for Firely Server specific operations. Currently just `'http://server.fire.ly/auth/scope/erase-operation'`
- **RequireClientSecret:** true / false - A *public client* cannot hold a secret, and then this can be set to `false`. Then the `ClientSecrets` section is ignored. See also the note below.
- **LaunchIds:** array of string - a 'launch id' could restrict access based on e.g. some currently active EHR context. Since Firely Auth is not connected to an EHR, this currently can only be set statically. Providing an empty array will make Firely Auth accept any launch id sent by the client (including none).
- **RefreshTokenLifetime:** If the client is allowed to use a *refresh token*, how long should it be valid? The value is in days. You can also use `HH:mm:ss` for lower values.

### Inferno test settings

The Inferno test suite for ONC Certification (g)(10) Standardized API has tests using the “Inferno-Public” client. For this client, `RequireClientSecret` has to be set to `false`. The same suite also issues a launch id as part of test 3.3. For this to succeed, add the launch id that is used to the `LaunchIds` array of the ‘Inferno’ client.

Below you will find the settings that can act as a reference for testing this suite. On top of that you will need to arrange:

For hosting (either directly with Kestrel as shown below, or with a reverse proxy that sits in front)

- SSL certificate for Firely Auth
- SSL certificate for Firely Server
- Configure both to use SSL protocols TLS 1.2 and 1.3

Necessary data:

- Pre-load one version of US-Core conformance resources to the Firely Server administration endpoint
- Pre-load the example resource of the same version of US-Core to the regular endpoint

We have a full walkthrough of Inferno testing available as a whitepaper, see [our resources](#).

### Firely Auth settings:

Put these settings in `appsettings.instance.json` next to the executable.

For Inferno you have to host it on https, with TLS 1.2 minimum. So you also need to provide a certificate for that (either to Kestrel as shown below, or to a reverse proxy that sits in front).

```
{
  "Kestrel": {
    "Endpoints": {
      "Http": {
        "Url": "http://localhost:5100"
      },
      "HttpsFromPem": {
        "Url": "https://localhost:5101",
        "SslProtocols": [ "Tls12", "Tls13" ],
        "Certificate": {
          "Path": "cert.pem",
          "KeyPath": "cert-key.pem"
        }
      }
    }
  },
  // Use "Https" option instead if you want to use a .pfx file. See https://docs.
  // microsoft.com/en-us/aspnet/core/fundamentals/servers/kestrel/endpoints
  },
  "FhirServerConfig": {
    "Name": "Firely Server",
    "FHIR_BASE_URL": "<url where you host Firely Server>"
  },
  "TokenConfig": {
    "AccessTokenType": {
      "Inferno": "Reference",

```

(continues on next page)



(continued from previous page)

```

    "Inferno-Public": "Reference"
  }
},
"KeyManagementConfig": {
  "RSA_Config": {
    "SupportedAlgorithms": [
      "RS256",
      "RS384",
      "RS512"
    ]
  },
  "EC_Config": {
    "SupportedAlgorithms": [
      "ES256",
      "ES384",
      "ES512"
    ]
  }
},
"TokenIntrospectionConfig": {
  "TokenIntrospectionResources": [{
    "Name": "Firely Server",
    "Secret": "secret"
  }]
},
"UserStore": {
  "Type": "InMemory", // InMemory | SqlServer
  "InMemory": {
    "AllowedUsers": [
      {
        "Username": "alice",
        "Password": "p@sSw0rd",
        "AdditionalClaims": [
          {
            "Name": "patient",
            "Value": "<id of a patient in your Firely Server, e.g. 'example'>"
          },
          {
            "Name": "fhirUser",
            "Value": "Practitioner/<id of a practitioner"
          }
        ]
      }
    ]
  }
},
"ClientRegistrationConfig": {
  "AllowedClients": [
    {
      "ClientId": "Inferno",
      "ClientName": "Inferno",
      "Enabled": true,

```

(continues on next page)

(continued from previous page)

```

    "RequireConsent": true,
    "RedirectUri": [ "https://inferno.healthit.gov/suites/custom/smart/launch",
↪ "https://inferno.healthit.gov/suites/custom/smart/redirect" ],
    "AllowedGrantTypes": [ "authorization_code" ],
    "ClientSecrets": [
      {
        "SecretType": "SharedSecret",
        "Secret": "secret"
      }
    ],
    "AllowFirelySpecialScopes": false,
    "AllowedSmartLegacyActions": [ "read", "write", "*" ],
    "AllowedSmartActions": [ "c", "r", "u", "d", "s" ],
    "AllowedSmartSubjects": [ "patient", "user" ],
    "AlwaysIncludeUserClaimsInIdToken": true,
    "RequirePkce": false,
    "AllowOfflineAccess": true,
    "AllowOnlineAccess": false,
    "RequireClientSecret": true,
    "LaunchIds": [ "<same id as for user>" ],
    "RefreshTokenLifetime": "90"
  },
  {
    "ClientId": "Inferno-Public",
    "ClientName": "InfernoPublic",
    "Enabled": true,
    "RequireConsent": true,
    "RedirectUri": [ "https://inferno.healthit.gov/suites/custom/smart/launch",
↪ "https://inferno.healthit.gov/suites/custom/smart/redirect" ],
    "AllowedGrantTypes": [ "authorization_code" ],
    "AllowFirelySpecialScopes": false,
    "AllowedSmartLegacyActions": [ "read", "write", "*" ],
    "AllowedSmartActions": [ "c", "r", "u", "d", "s" ],
    "AllowedSmartSubjects": [ "patient", "user" ],
    "AlwaysIncludeUserClaimsInIdToken": true,
    "RequirePkce": false,
    "AllowOfflineAccess": true,
    "AllowOnlineAccess": false,
    "LaunchIds": [],
    "RequireClientSecret": false,
    "RefreshTokenLifetime": "90"
  },
  {
    "ClientId": "Inferno-Bulk",
    "ClientName": "InfernoBulk",
    "Enabled": true,
    "RedirectUri": [ "https://inferno.healthit.gov/suites/custom/smart/launch",
↪ "https://inferno.healthit.gov/suites/custom/smart/redirect" ],
    "AllowedGrantTypes": [ "authorization_code", "client_credentials" ],
    "AllowFirelySpecialScopes": false,
    "AllowedSmartLegacyActions": [ "read" ],
    "AllowedSmartActions": [ "c", "r", "u", "d", "s" ],

```

(continues on next page)

(continued from previous page)

```

    "AllowedSmartSubjects": [ "system" ],
    "RequirePkce": false,
    "AllowOfflineAccess": true,
    "AllowOnlineAccess": false,
    "ClientSecrets": [
      {
        "SecretType": "JWK",
        "Secret": "{ 'e': 'AQAB', 'kid': 'b41528b6f37a9500edb8a905a595bdd7', 'kty': 'RSA',
→ 'n': 'vjbIzTqiY8K8zApeNng5ekNNIxJfXAue9BjoMrZ9Qy9m7yIA-
→ tf6muEupEXWhq70tC7vIGLqJJ408m7yiH8H2qklX2mCAMg3xG3nbykY2X7JXtW9P8VIdG0sAMt5aZQnUGCgSS3n0qaoGn2LUlTGI
→ 4Nr9_3Ki3UCiICeCiAE224jGCg001QU6qj2gEB3o-DWJF1G_dz1y-Mxo5ivaeM0vWuodjDrp-aiabJcSF_
→ dx26sdC9dZdBKXFDq0t19I9S9AyGpGDJwzGRtWHY6LsskNHLvo8Zb5AsJ9eRZKpnh30SYBZI9WHtzU85M9WQqdScR69Vyp-
→ 6Uhfbvw' }"
      },
      {
        "SecretType": "JWK",
        "Secret": "{ 'kty': 'EC', 'crv': 'P-384', 'x':
→ 'JQKTsV6PT5Szf4QtDA1qrs0EJ1pbimQmM2SKvz0lIAqlph3h10HmZ2i7MXahIF2C', 'y':
→ 'bRWWQRJBgDa6CTgwofYrHjVGc0-A7WNEnu4oJA5OUJPPPpczgx1g2NsfinK-D2Rw', 'key_ops': ['verify
→ '], 'ext': true, 'kid': '4b49a739d1eb115b3225f4cf9beb6d1b', 'alg': 'ES384' }"
      }
    ],
    "RequireClientSecret": true,
    "RefreshTokenLifetime": "90"
  }
}

```

## Firely Server settings

Put these settings in appsettings.instance.json, next to the executable.

For Inferno you have to host it on https, with TLS 1.2 minimum. So you also need to provide a certificate for that (either to Kestrel as shown below, or to a reverse proxy that sits in front).

```

"Hosting": {
  "HttpPort": 4080,
  "HttpsPort": 4081, // Enable this to use https
  "CertificateFile": "<your-certificate-file>.pfx", //Relevant when HttpsPort is present
  "CertificatePassword" : "<cert-pass>", // Relevant when HttpsPort is present
  "SslProtocols": [ "Tls12", "Tls13" ] // Relevant when HttpsPort is present.
},
"SmartAuthorizationOptions": {
  "Enabled": true,
  "Filters": [
    {
      "FilterType": "Patient",
      "FilterArgument": "_id=#patient#"
    }
  ],

```

(continues on next page)

(continued from previous page)

```

"Authority": "<url where Firely Auth is hosted>",
"Audience": "Firely Server",
"RequireHttpsToProvider": true,
"Protected": {
  "InstanceLevelInteractions": "read, vread, update, patch, delete, history, ↵
↵conditional_delete, conditional_update, $validate, $meta, $meta-add, $meta-delete,
↵$export, $everything, $erase",
  "TypeLevelInteractions": "create, search, history, conditional_create, compartment_
↵type_search, $export, $lastn, $docref",
  "WholeSystemInteractions": "batch, transaction, history, search, compartment_system_
↵search, $export, $exportstatus, $exportfilerequest"
},
"TokenIntrospection": {
  "ClientId": "Firely Server",
  "ClientSecret": "secret"
},
"ShowAuthorizationPII": false,
// "AccessTokenScopeReplace": "-",
"SmartCapabilities": [
  "LaunchStandalone",
  "LaunchEhr",
  // "AuthorizePost",
  "ClientPublic",
  "ClientConfidentialSymmetric",
  // "ClientConfidentialAsymmetric",
  "SsoOpenidConnect",
  "ContextStandalonePatient",
  "ContextStandaloneEncounter",
  "ContextEhrPatient",
  "ContextEhrEncounter",
  "PermissionPatient",
  "PermissionUser",
  "PermissionOffline",
  "PermissionOnline",
  "PermissionV1",
  // "PermissionV2",
  "ContextStyle",
  "ContextBanner"
]
},
// PipelineOptions: make sure that Vonk.Plugin.SoFv2 is enabled
"PipelineOptions": {
  "PluginDirectory": "./plugins",
  "Branches": [
    {
      "Path": "/",
      "Include": [
        // all other default plugins...
        "Vonk.Plugin.SoFv2",
      ],
      "Exclude": [
        // ...

```

(continues on next page)

(continued from previous page)

```
    ],
    {
      "Path": "/administration",
      "Include": [
        //...
      ],
      "Exclude": [
        //...
      ]
    }
  ]
}
```

## 12.2.4 Deployment of Firely Auth

### License

You can use your Firely Server license file, provided that it contains the license token for Firely Auth: `http://fire.ly/server/auth`. This token is included in the evaluation license for Firely Server, and in your production license if Firely Auth is included in your order.

Configure the path to the license file in the appsettings, section *License*.

### Executable / binaries

You can request a copy of the binaries from Firely through our *contact page*

### Docker image

A docker image is available on the Docker hub, under *firely/auth*.

See the instructions on *running Firely Server in Docker* to learn about adjusting settings and providing the license file. Firely Auth is configured in the same way.

### InMemory user store

The InMemory user store is only meant for testing your setup or evaluating Firely Auth. For production use configure the SQL Server user store.

The users for the InMemory user store can be configured in *User store*

### SQL Server user store

Use of the SQL Server user store requires Microsoft SQL Server version 2016 or newer.

Using your favourite database administration tool:

- create a new database, e.g. 'firely\_auth\_store'
- in this database, execute the script `scripts/InitializeSchema.sql`, available in the binaries
- create a connection string to this database
- configure *User store*

```
{
  "Type": "SqlServer",
  "SqlServer": {
    "ConnectionString": "<connectionstring from previous step>"
  }
}
```

In the connection string you can use a user that is only allowed to read and write from the existing tables, no further DDL is needed.

To add users to the store, you can use the *Firely Auth User Management*.

### 12.2.5 Glossary of terms in authorization

#### user

The patient, practitioner or other person that uses a *client* to access resources on Firely Server. In OAuth terms this is the 'resource owner'.

#### client

An app or website that accesses resources on Firely Server, either on behalf of a *user* or directly.

#### fhir server

An instance of Firely Server hosting resources that can be accessed by a *client* on behalf of a *user*. In OAuth terms this is usually referred to as a 'resource server'.

#### authorization server

An instance of Firely Auth that can provide a *access token* and potentially *refresh token* to a *client*

#### access token

A string issued by the *authorization server* to the requesting *client*, that the client then needs to present to the *fhir server* to get access to resources. The fhir server will check the validity of the token. The client need not understand the token or its format and is not supposed to read it.

An access token issued by Firely Auth can be in one of two forms: a *JWT*, or a *reference token*.

#### JWT

JWT is an abbreviation for JSON Web Token. It defines a structure for an access token, based on JSON, with a predefined list of elements. A JWT is self contained, so it has all the data in it for the *fhir server* to check its validity. This however means that the token cannot be revoked in case of malicious use. The token remains valid until its expiration time is past.

A JWT is encoded into a single string. You can easily decode it on [JWT.io](https://jwt.io).

#### JWK

JWK is an abbreviation of JSON Web Key. It defines a JSON structure for a key. A client can sign requests with a private key and publish a JWK for the corresponding public key for Firely Auth to use when validating the

signature. Multiple JWKs are usually published on a url in the form of a JWKS - a JSON Web Key Set. Multiple keys allow for key rotation without downtime.

#### reference token

A reference token is a string that does not contain any information about what is granted. Instead it is a *reference* to the actual token that is kept within Firely Auth. For Firely Server to check a reference token it sends the token to the *token introspection endpoint* of Firely Auth. Firely Auth will then send the validity along with relevant claims to Firely Server. A reference token is slower to validate than a *JWT* because of the roundtrip to Firely Auth. But it is safer than a *JWT* because the *client* cannot read any information from the token and the token can be revoked before its original expiration time.

#### authorization code

A code (a string) returned by Firely Auth to a *client* as part of an *authorization code flow*. The client can then exchange the authorization code for an *access token*.

#### authorization code flow

A flow to authorize a *client* to request an *access token* from Firely Auth. Typically used for a *public client*. See <https://oauth.net/2/grant-types/authorization-code/>. It can be further secured by the use of *PKCE*.

#### PKCE

Abbreviation of Proof Key for Code Exchange. A protocol used in conjunction with *authorization code flow* to prevent certain attacks. It uses temporary secrets generated by the *client*.

#### grant type

A way of authorizing a *client* to request an *access token* from Firely Auth. Also referred to as authorization flows. Firely Auth supports two grant types:

- *client credentials*
- *authorization code flow* (optionally with *PKCE*)

#### client credentials

A *confidential client* may use credentials (similar to username / password, but for software) to request an *access token* directly from Firely Auth.

#### claim

An element in an *access token* that tells Firely Server that the *client* claims to have access to certain resources. A client requests a claim. Then Firely Auth will optionally ask the *user* for consent to grant this claim to the client. Then Firely Auth includes the claim in the access token. The client presents the access token as part of its request to Firely Server. Finally Firely Server must understand the claim and will restrict access to resources accordingly.

Firely Server and Firely Auth understand the *claims as defined by SMART on FHIR* (v1 and v2).

#### refresh token

The Refresh Token grant type is used by clients to exchange a refresh token for an access token when the access token has expired. This allows clients to continue to have a valid access token without further interaction with the user.

#### token introspection endpoint

Endpoint offered by Firely Auth to inspect a *reference token*.

#### confidential client

A *client* that can guard a shared secret. E.g. a web application, where the code lives on a well managed server. See also *SMART on FHIR clients*.

#### public client

A *client* that cannot guard a shared secret, e.g. because its code is public. Like in native apps or some Single Page Applications. See also *SMART on FHIR clients*.

## 12.2.6 Firely Auth User Management

When using the SQL Server user store, you can manage the users therein with a command line utility called the ‘Firely Auth Management app’. This can be used both interactively from the commandline and automated from a script or devops pipeline.

### Executable

The app is included in the *binaries* of Firely Auth. You can invoke it with:

```
./Firely.IdentityServer.ManagementApp.exe
```

### Commands and parameters

For brevity here is a full list of the available commands and their parameters.

The parameter ‘`--connectionString`’ or `-cs` is used multiple times. Use the same connection-string as in *SQL Server user store*. You can set the connectionstring in an environment variable `FIRELY_IDENTITY_SERVER_USERS_CONNECTION_STRING` so you don’t need to repeat it all the time.

- `./Firely.IdentityServer.ManagementApp.exe`
  - –?, –h, –help: Show available commands
  - users: invoke user management
    - list: list all users
      - cs, –connectionString
    - create: create a new user, specifying username and password
      - cs, –connectionString –u, –username: username –p, –password: password - enclose in double quotes if it contains spaces. –cl`, –claim <name=value> [<name=value> ...]: list of additional claims for this user, see ‘AdditionalClaims’ under User store
    - delete: delete a user so the user no longer has access
      - u, –username: username –cs, –connectionString`

### Password requirements

The password for a new user must conform to all of these requirements:

- at least 6 characters length
- at least 1 lowercase character (a - z)
- at least 1 non-alphanumeric characters
- at least 1 uppercase character (A - Z)



## Examples

For all examples you need to provide a connectionString. It is easiest to set this upfront in an environment variable:

```
$env:FIRELY_IDENTITY_SERVER_USERS_CONNECTION_STRING = "Server=localhost,1433;
↳MultipleActiveResultSets=true;Database=firely_auth_store;User Id=<db_user>;Password=
↳<db_user_password>;Encrypt=True"
```

### List the users in the database

```
./Firely.IdentityServer.ManagementApp.exe users list
```

### Create a new user in the database

```
./Firely.IdentityServer.ManagementApp.exe users create -u bob -p 1P@ssword -c_
↳patient=Patient/bob
```

Try to issue the same command again to see that a user with the same name is rejected.

### Delete a user from the database

```
./Firely.IdentityServer.ManagementApp.exe users delete -u bob
```

### Change a user's password

There is no function to change a user's password directly. Instead delete the user and recreate it with the new password and any claims.

## 12.2.7 Firely Auth Bill of Materials

Firely Auth is mainly built using libraries from Duende, Microsoft .Net Core and ASP.NET Core, along with a limited list of other libraries. This is the full list of direct dependencies that Firely Auth has on other libraries, along with their licenses.

This list uses the NuGet package names (or prefixes of them) so you can easily lookup further details of those packages on [NuGet.org](https://www.nuget.org) if needed.

1. Duende.IdentityServer.\*: <https://duendesoftware.com/license>
2. Microsoft.AspNetCore.\* - Apache 2.0
3. Microsoft.EntityFrameworkCore.\* - MIT
4. Microsoft.Extensions.Configuration: MIT
5. System.Commandline.\*: MIT
6. Serilog(\*) - Apache-2.0
7. HL7.Fhir.\* - Firely OSS license (see *Firely OSS License*)

- 8. Simplifier.Licensing - as Hl7.Fhir
- 9. CreativeCode.JWK - MIT
- 10. Newtonsoft.Json - MIT
- 11. Superpower - Apache 2.0

Firely Server is capable of generating Swagger / OpenAPI documentation. The content of these definitional artifacts are based on a [CapabilityStatement](#). A full overview of the REST API provided by Firely Server can be found here: [OpenAPI documentation](#). Please note that due to the feature-richness of the CapabilityStatement resource, it is not possible to expose all information through the OpenAPI documents and some limitations exist.

## 13.1 Limitations of Capability Statement to OpenAPI mapping

- Resource profile definition references are not displayed.
- Field profile definition references are not displayed.
- Supported interactions are not displayed.
  - i.e. `searchInclude`, `searchRevInclude`
- Supported interactions are not directly listed, but are inferred from the output.
  - i.e. `GET {BASE_URL}/Appointment/{logical_id}/_history` infers `readhistory` is permitted.

## 13.2 Generating OpenAPI for a local instance of Firely Server

OpenAPI documents can be generated for a local instance of Firely Server using the open source github repository [Microsoft/Fhir-CodeGen](#).

A sample configuration is listed below:

```
fhir-codegen-cli
  --fhir-server-url <FhirServerUrl>
  --resolve-external false
  --language-options
  → "OpenApi|SingleResponses=false|Metadata=true|SchemaLevel=names|MultiFile=true"
  --output-path <OpenApiOutputDirectory>
  --language OpenApi
```

See Microsoft's documentation for other configuration options and usage details.



## **CONTACT US**

For questions, feedback, or consultancy for Firely Server, please send an e-mail to [server@fire.ly](mailto:server@fire.ly). We look forward to hearing from you!



## A

access token, **344**  
 Add()  
     built-in function, 293, 296  
 AddIf()  
     built-in function, 293, 296  
 AddIfNotExists()  
     built-in function, 293, 294, 296  
 AddOrReplace()  
     built-in function, 294  
 AddParent()  
     built-in function, 296  
 AddTreePath()  
     built-in function, 296  
 AndInformationModel()  
     built-in function, 303  
 AndInteraction()  
     built-in function, 303  
 AndMethod()  
     built-in function, 303  
 AndResourceTypes()  
     built-in function, 303  
 AndStatusCodes()  
     built-in function, 303  
 AnnotateWithSourceNode()  
     built-in function, 295  
 authorization code, **345**  
 authorization code flow, **345**  
 authorization server, **344**

## B

built-in function  
     Add(), 293, 296  
     AddIf(), 293, 296  
     AddIfNotExists(), 293, 294, 296  
     AddOrReplace(), 294  
     AddParent(), 296  
     AddTreePath(), 296  
     AndInformationModel(), 303  
     AndInteraction(), 303  
     AndMethod(), 303  
     AndResourceTypes(), 303

AndStatusCodes(), 303  
 AnnotateWithSourceNode(), 295  
 Cache(), 296  
 Child(), 294, 296  
 ChildString(), 294, 296  
 DefinitionSummary(), 296  
 ForceAdd(), 294  
 ForcePatch(), 295  
 ForcePatchAt(), 295  
 GetParent(), 296  
 GetTreePath(), 296  
 OnCustomInteraction(), 303  
 OnInteraction(), 303  
 Patch(), 295  
 Relocate(), 295  
 Remove(), 295  
 RemoveEmptyNodes(), 294  
 Rename(), 295  
 Revalue(), 295  
 SelectNodes(), 295  
 SelectText(), 295

## C

Cache()  
     built-in function, 296  
 Child()  
     built-in function, 294, 296  
 ChildString()  
     built-in function, 294, 296  
 claim, **345**  
 client, **344**  
 client credentials, **345**  
 confidential client, **345**

## D

DefinitionSummary()  
     built-in function, 296

## F

fhir server, **344**  
 ForceAdd()  
     built-in function, 294

ForcePatch()  
    built-in function, 295  
ForcePatchAt()  
    built-in function, 295

## G

GetParent()  
    built-in function, 296  
GetTreePath()  
    built-in function, 296  
grant type, 345

## J

JWK, 344  
JWT, 344

## O

OnCustomInteraction()  
    built-in function, 303  
OnInteraction()  
    built-in function, 303

## P

Patch()  
    built-in function, 295  
PKCE, 345  
public client, 345

## R

reference token, 345  
refresh token, 345  
Relocate()  
    built-in function, 295  
Remove()  
    built-in function, 295  
RemoveEmptyNodes()  
    built-in function, 294  
Rename()  
    built-in function, 295  
Revalue()  
    built-in function, 295

## S

SelectNodes()  
    built-in function, 295  
SelectText()  
    built-in function, 295

## T

token introspection endpoint, 345

## U

user, 344