

High-performance Storing of Conformance Data

Ewout Kramer, CTO Firely



HL7 FHIR DevDays International 2022 | Hybrid Edition, Cleveland, OH | June 6–9, 2022 | @HL7 | @FirelyTeam | #fhirdevdays | www.devdays.com

ORGANIZED BY

firely

HL7[®]
International

Who am I?

Ewout Kramer

CTO @ Firely

.NET FHIR SDK developer

Technology stack & architecture
@ Firely



Most of our tools need FHIR metadata (aka Conformance resources)



How do we distribute this data?

How do we access this data quickly?

SIMPLIFIER.NET Search

Packages Search
This registry contains all published FHIR packages which can be downloaded or referenced from your projects.

Search packages All FHIR Versions


de.gematik.isik-basismodul-stufe1 1.0.4	FHIR R4	created 34 minutes ago
nictiz.fhir.nl.stu3.questionnaires 2.0.7 Nictiz NL package of FHIR STU3 conformance resources for the Questi...	FHIR STU3	created about 4 hours ago
nictiz.fhir.nl.stu3.zib2017 2.2.0 Nictiz NL package of FHIR STU3 conformance resources for HCIM Rele...	FHIR STU3	created about 5 hours ago
iteyes.myhw.core 1.0.9 ObservationLab -> Observation ObservationFunc -> Observation by jh...	FHIR R4	created about 6 hours ago
ca.on.oh.mha-pds 0.9.0-alpha-1.9.7.4 testing MRN identifier system with binding	FHIR R4	created about 17 hours ago
uk.nhsdigital.r4.test 2.6.8-prerelease	FHIR R4	created about 21 hours ago
uk.nhsdigital.medicines.r4.test 2.6.2-prerelease	FHIR R4	created about 23 hours ago
hl7.fhir.au.base 4.0.0	FHIR R4	created a day ago

The packaging is solved, yes, but....

- We find ourselves unzipping (tarballing) it into cached directories
- Putting the contents in databases
- Parsing Json/XML for every access
- Unresolved references to ValueSets, StructureDefinitions (@runtime!)
- Preprocessing/cleaning
- Need the closure of the packages, a complete set
- Horrible cold-start times, waiting for cache build-up...

Solution: Conformance Archives (aka “CAR files”)

“NPM files on Speed”



- Hundreds of huge xml/json files
- Packaged in NPM archives.
- Need to parse & index
- Would be a bottleneck if stored in a shared database

Readonly and static once deployed


“IG in a file”

Calculate Closure of dependencies

StructureDefinition →
Snappen →
ElementSchema

ValueSets → Expansion →
Lookup tables

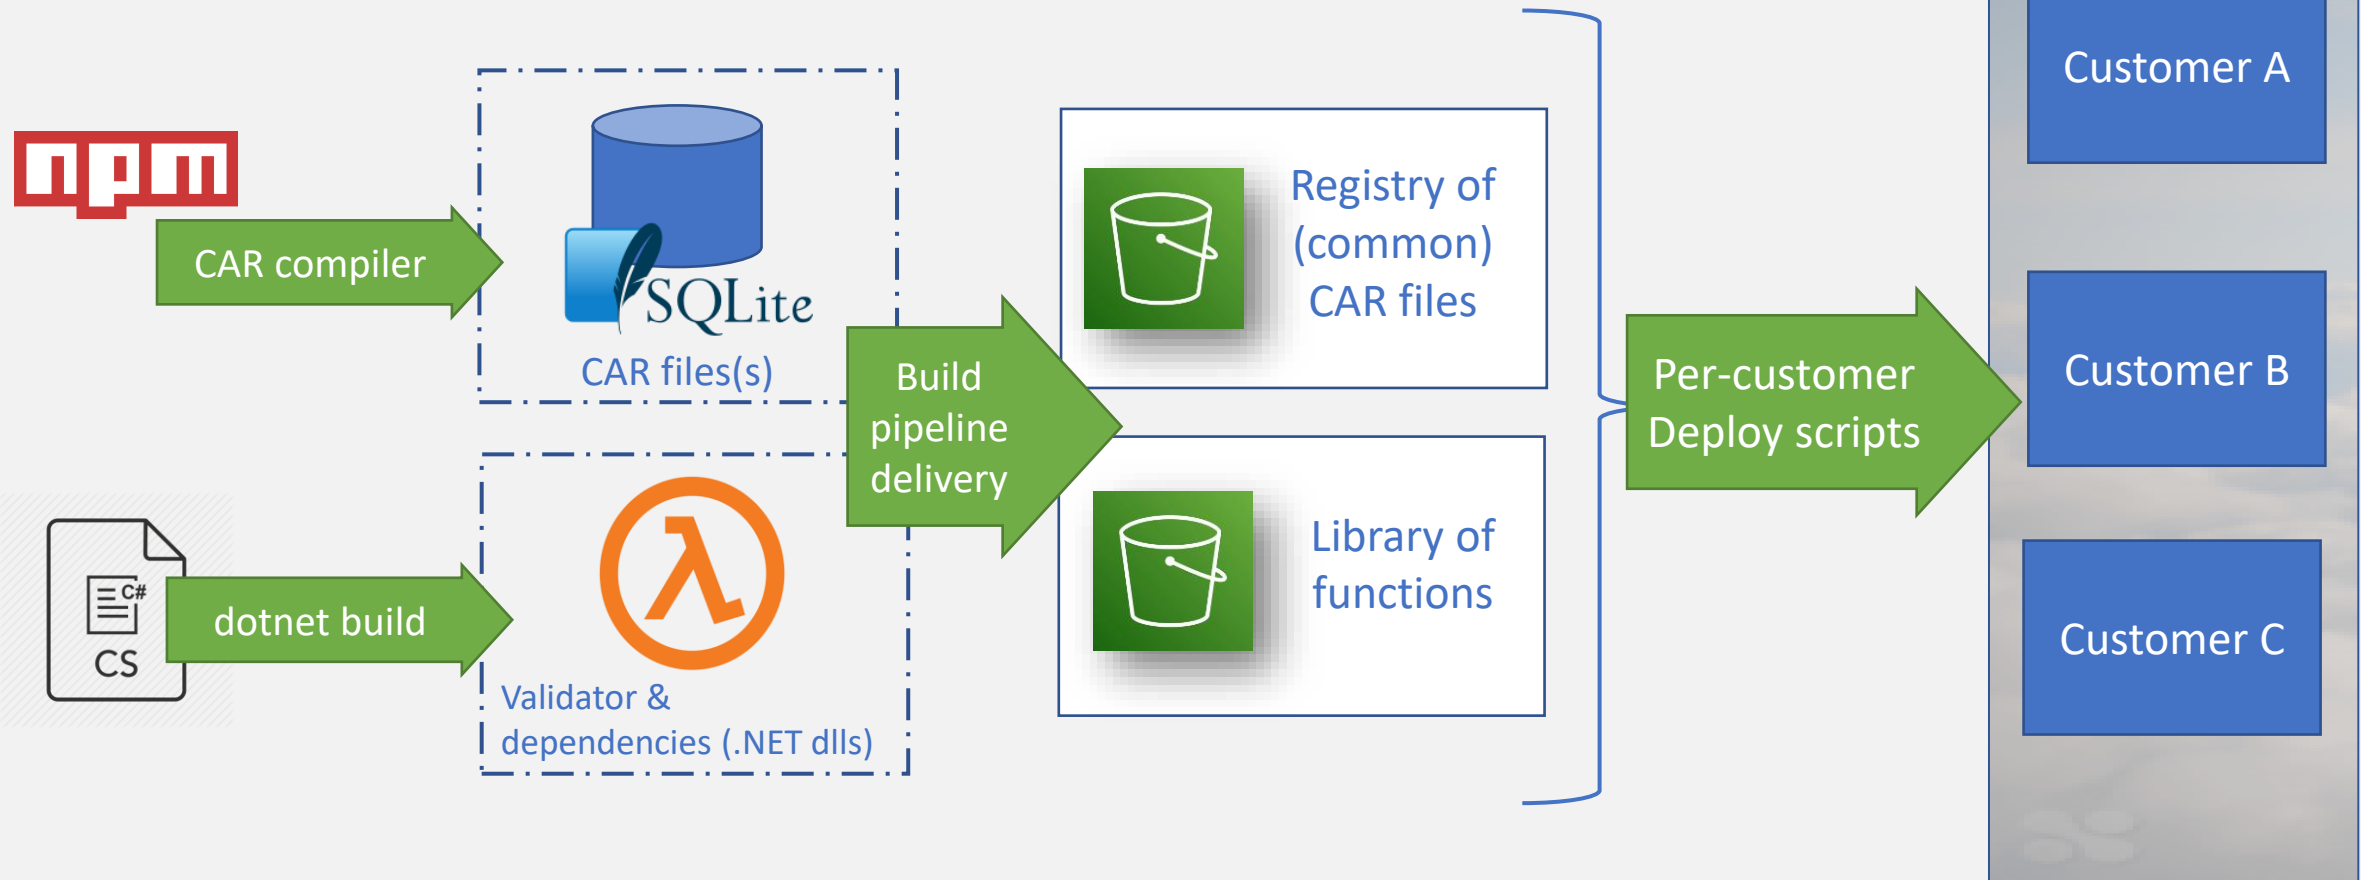
SearchParameters →
Denormalized structures



SQLite

“optimal (verified & denormalized) binary format”

Build, Delivery & Deployment



- Does this sound familiar?
- What are your solutions that you can share?

Our own exciting ideas :

- Maybe build a Filesystem-with-metadata-on-Sqlite?
- Filesystem in Userspace (FUSE)?

ORGANIZED BY

