


From Zero to Production FHIR Server in 6 Months at a Large Health System

Nic Lorenzen Director of Software Engineering



HL7 FHIR DevDays 2021, Virtual Edition, June 7–10, 2021 | @HL7 | @FirelyTeam | #fhirdevdays | www.devdays.com

ORGANIZED BY



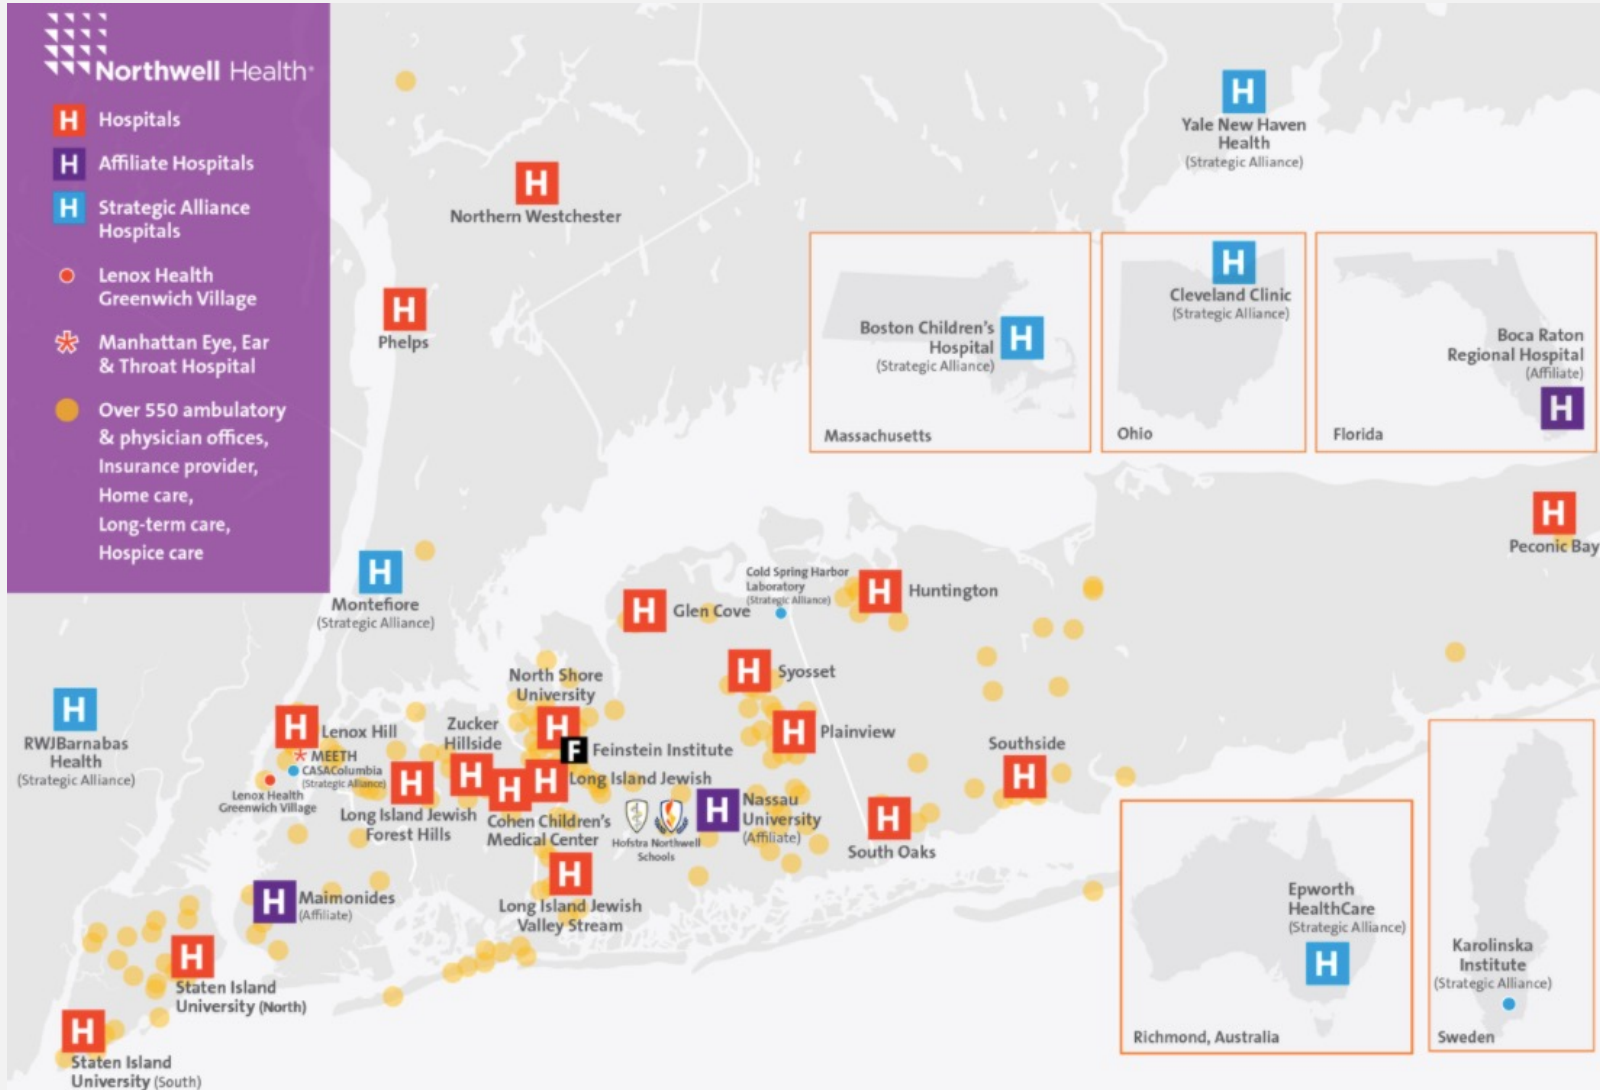
PARTNER



Who am I?

- Nicholas (Nic) Lorenzen
- Director of Software Engineering
- Northwell Emerging Technology and Innovation Team at Northwell Health
- Enjoys hands-on development and architecture of cloud native applications in the healthcare domain.

What makes Northwell Health a large health system?



- 14th largest health system in the U.S. and largest in New York
- 21 Hospitals, 6,600 hospital and long-term beds
- 61,000+ employees
- 550+ Ambulatory and physician practices

FHIR at Northwell:



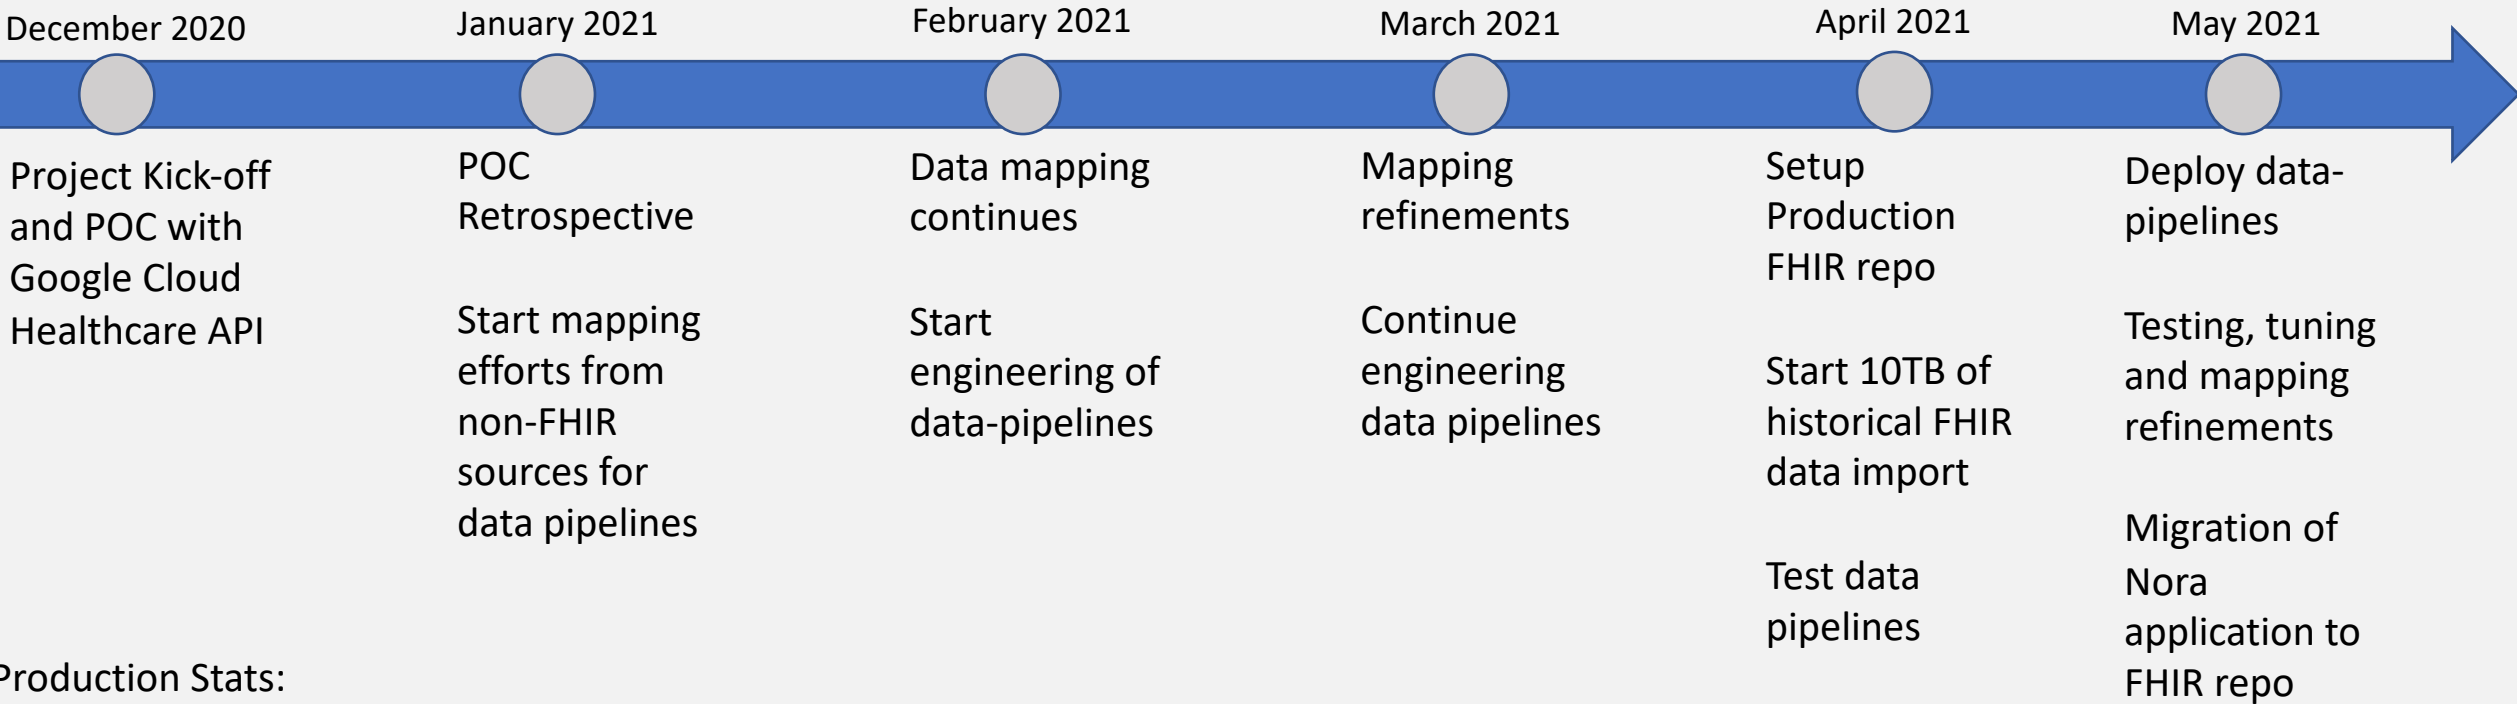
For more information on Nora: <https://northwell.edu/nora>

Learning from the past and moving forward:

Past FHIR repo attempts left us with a lot of expensive infrastructure and unacceptable performance for production use. Going forward we wanted:

- Fully managed cloud based FHIR repo.
 - Turn the key and go with no knobs or dials to turn for more storage and compute.
- Cost effective and predictable billing.
 - Pay for what you use and nothing more.
- High performance, even with complex searches.
 - Our users expect sub 1 second response times or better.
- Support for most of the search capabilities defined in the spec (i.e., `_lastn`, `_include`, `_revinclude`, `chaining`, etc.).

Our Journey:



Production Stats:

- 2 million+ HL7 derived bundles (10-15 million individual resources) posted to FHIR repo per day
 - 20-75 bundles received per/sec
- Easily supports 15k+ Nora users
- Stores 200-500 GB of data per month

Lessons learned:

- Mapping never finishes, even when you're done.
 - You're never going to get it right the first time; have processes prepared to crawl through resources for data corrections.
- Be elegant when integrating data sources to FHIR.
 - Understand the events and only insert/update what's necessary when it's necessary.
 - Example: update insert Patient and Encounter with registration (ADT A01), but not with an order (ORM O01). These rules may vary depending on your org.
- Prioritize patient merges.
 - Patient merges in FHIR is far from mature. Figure out a process that works for your org, develop it and test it asap.
- Parallelization + bundles is your (difficult) friend.
 - Use bundles to batch resources when possible. It will save money and improve performance.
 - Analyze your data and you'll find points where you can still maintain ordering with parallelization. It may result in non-traditional ingestion steps, but you'll gain performance and potentially save money.

Contact

- During DevDays, you can find / reach me here:
 - Via Whova App – Speaker's Gallery
 - Email: nlorenzen@northwell.edu

Q&A

ORGANIZED BY



PARTNER

