Optimizing Applications for Expedited Risk Assessment on Amazon Web Services (AWS) at the FDA

Hong Yang[1], Yin Huang[1], Rebecca Kahn[2] Mark Walderhaug[1], Jason Claeys[1], Richard Forshee[1]

1. FDA/CBER; 2. SAIC

Abstract

Background

A major challenge for the FDA Center for Biologics Evaluation and Research (CBER) is evaluating and responding quickly to the risk of emerging transfusion-transmitted diseases that may impact the blood supply. When a new transfusion-transmitted disease emerges, detailed information about the disease is often limited, scattered across myriad sources, which complicates risk evaluation for regulatory decision-making.

Methods

CBER's Office of Biostatistics and Pharmacovigilance (OBPV) created two decision support tools: GIS-based Risk Evaluation and Assessment Tool (GREAT) web application, which provides geographic risk information, and Blood Risk (BRisk) Tool, which performs risk assessment using probabilistic algorithms in predetermined models. These tools were developed as web applications and have been deployed to the FDA's Amazon Web Services (AWS) GovCloud West organization. Most recently, AWS Step Functions were incorporated to update S3 disease caches, which are used to speed up response time since Lambda functions including Java code have a long start up time.

Results

The applications help users evaluate risk and evaluate potential mitigation policies. Speed was improved by pulling data from Step Function created caches.

Conclusions/Implications

The AWS-based tools provide decision-makers with information concerning identified threats and possible mitigations and assist in the development and evaluation of emerging donor deferral and blood screening policies. These applications utilize AWS to take advantage of the flexibility and scalability the cloud provides.

Introduction

The BRisk and GREAT applications were developed to expedite the process of evaluating potential policy responses to emerging or resurging infectious diseases that may spread through the blood supply.

First, GREAT aggregates geographic data that is spread across a variety of sources and displays this information on a map. GREAT provides quick estimates of the impact of potential mitigation policies. Once the regions that contribute the most risk have been identified, BRisk can be used to conduct a more detailed analysis of the impact of potential mitigation policies, which could include conducting blood testing or deferring at-risk donors from donating blood.

GREAT

GREAT assists modelers and decision-makers in evaluating the latest public health data from the perspective of geographic risk. The tool provides a rapid overview of changes in emerging infectious diseases using data from a variety of sources and identifies the regions with the most risk (e.g. ranked by risk contribution that takes into account the potential number of U.S. donors exposed in the region) for further analysis.

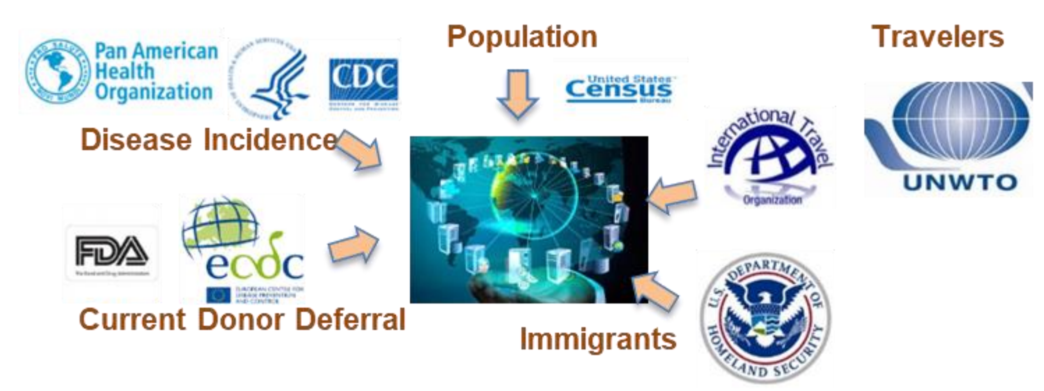

Figure 1: GREAT incorporates data from a variety of online sources

GREAT's homepage includes an interactive map. Users can select which data attributes should be displayed on this map, including disease, traveler, population, and immigration data as well as risk contribution and donor loss calculations. These risk contribution and donor loss calculations take into account any mitigation policies configured by the user. GREAT also generates charts of disease trends over time.

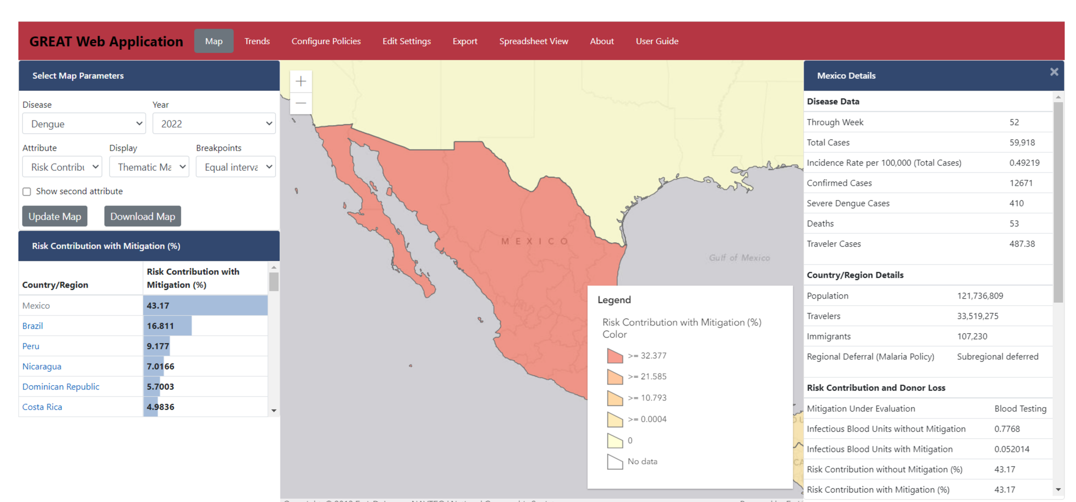

Figure 2: GREAT's homepage features an interactive map

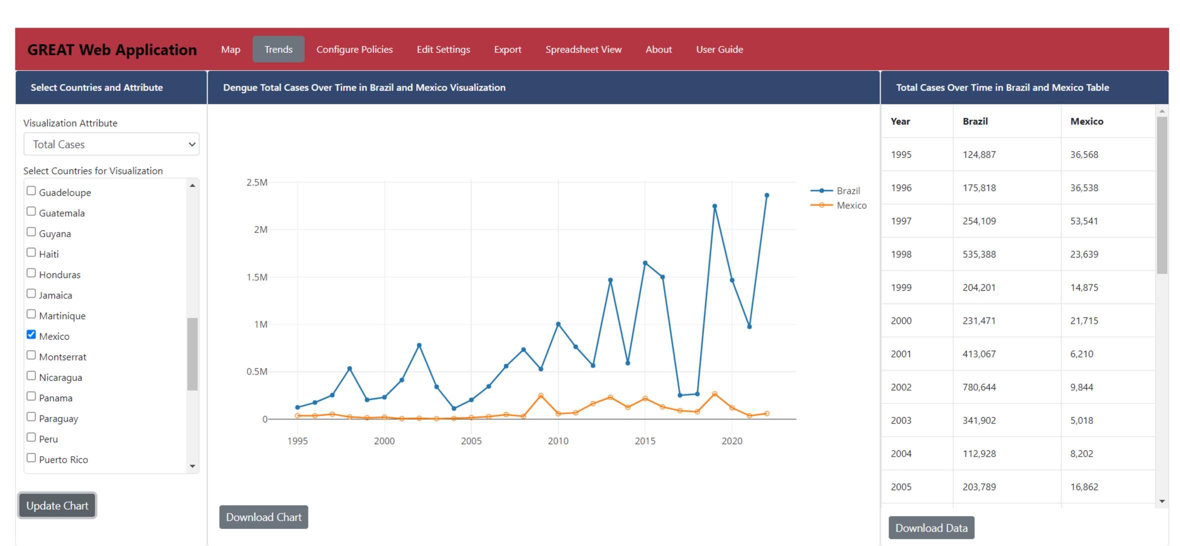

Figure 3: GREAT's homepage features an interactive map

BRisk

Once GREAT has been used to identify regions of interest and potential intervention policies, BRisk is used to conduct more detailed risk assessment using probabilistic algorithms in predetermined models. Three model templates, which model different exposure mechanisms (travel, behavior, or residence), are used to create models for emerging infectious diseases using Lumina's Analytica. These templates are designed to be modular to be easily modified by users. They contain thousands of interconnected nodes but only a small subset of these nodes need to be changed to model a new disease.

The model templates consist of three modules:

Module 1: Disease prevalence estimates the number of cases of the disease in the population.

Module 2: Infectious blood units calculates the number of infectious blood units, taking into account existing and alternative risk mitigation strategies, such as donor deferral and/or blood testing.

Module 3: Transfusion-transmitted risk determines the number of infections using dose-response information.

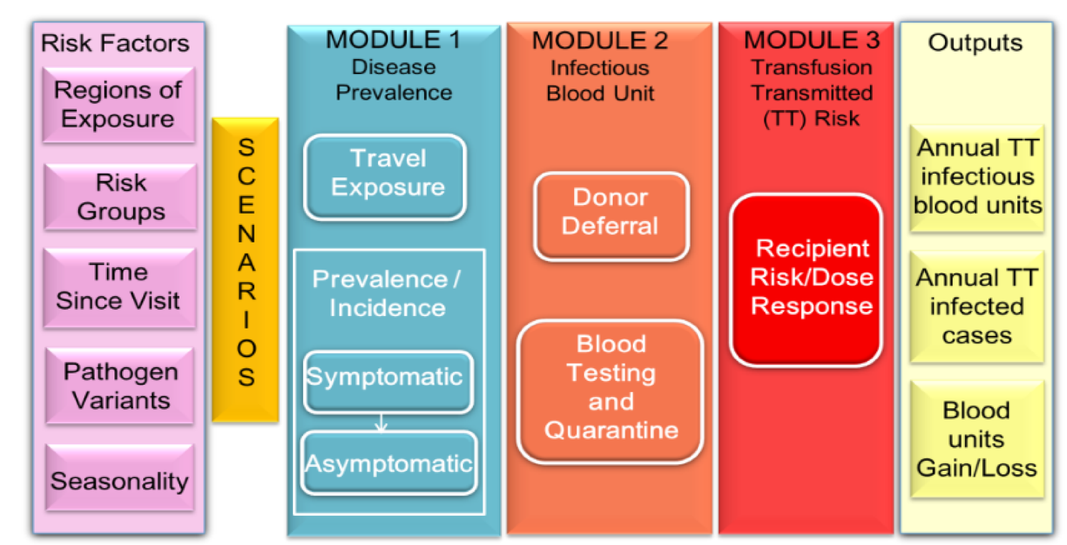

Figure 4: BRisk model structure

BRisk assists users in updating these templates so that they model an emerging disease by only exposing the nodes that need to be updated within an easy-to-use web user interface.

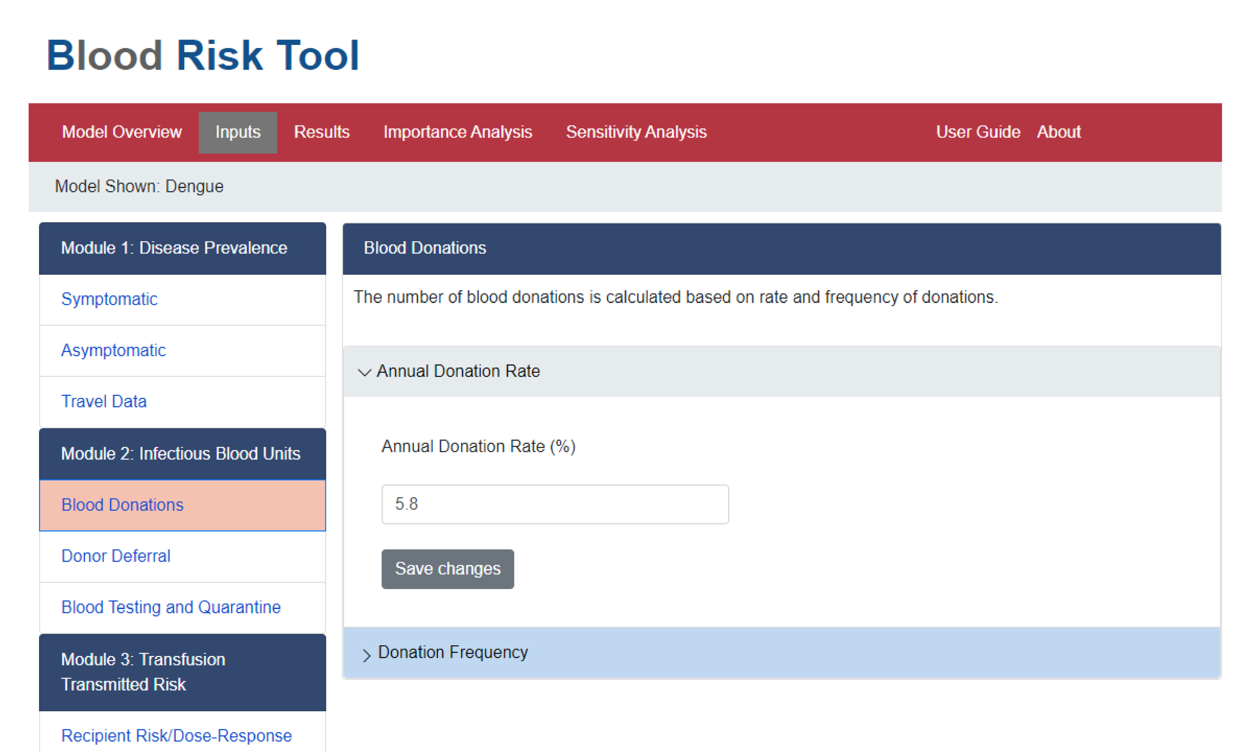

Figure 5: Modifying model inputs in BRisk

Lessons Learned

When architecting an application for AWS, utilizing Lambda functions can be an attractive option. The primary benefit of Lambda functions is that they only run when requested, unlike servers, which need to run all of the time to respond to requests. Avoiding having a server run all the time can represent cost savings.

However, for GREAT, Lambda functions did not represent the ideal architecture. The Lambda functions required enough time to load prerequisites and dependencies each time a function was called that speed was an issue. To resolve this issue, Step functions and S3 storage cache were added to the architecture. Step functions can be used to orchestrate Lambda functions and were used to call Lambda functions and save the results to S3 storage. Then new Lambda functions were created to fetch data from the S3 storage and return it to users. These changes caused dramatic speed improvements. The lesson learned is that using serverless Lambda functions is not always the most efficient option.

Conclusion

Developers are deploying the applications for production use by selected FDA Office of Blood Research and Review users on the FDA GovCloud West Organization. These applications could be used to conduct other types of risk assessment within CBER and at other centers at FDA.