FAMINE EARLY WARNING SYSTEMS NETWORK
PILLAR 2: LEARNING AND DATA HUB

Annual Plan for Revisions to the FEWS NET Data Management Platform

Volume 2. Methodology and Findings from Technical and Stakeholder Analyses

The Hub Task Order 1

Submission Date: August 14, 2020

This supplement is provided as an addition to the *Annual Plan for Revisions to Data Management Platform for the FEWS NET Hub*'s second year of the project. This supplement was referenced within the text and has been made available here for more in-depth information on the topic.

CONTENTS

ACRONYMS AND ABBREVIATIONS

ACLED	Armed Conflict Location & Event Data
ADS	Automated Directive Service
AMIS	Agricultural Management Information System
API	Application Programming Interface
ATO	Authority to Operate
AWS	Amazon Web Services
BHA	Bureau for Humanitarian Affairs
BRFS	Bureau for Resilience and Food Security
CIO	Chief Innovation Office
CKAN	Comprehensive Knowledge Archive Network
COR	Contracting Officer's Representative
CRM	Customer Relationship Manager
CSAM	Cyber Security Assessment and Management
CSV	Comma-Separated Values
CPI	Consumer Price Index
DDATA	Disaster Data, Assessments, Technology & Analysis (of USAID/BRFS)
DEC	Development Experience Clearinghouse
DMP	Data Management Platform
DRF	Django Rest Framework
ECS	Elastic Container Service
EKS	Elastic Kubernetes Service
ETL	Extract, Transform, Load
EW	Early Warning
FAO	Food and Agriculture Organization
FDE	FEWS NET Data Explorer
FDW	FEWS NET Data Warehouse
FedRAMP	Federal Risk and Authorization Management Program
FEWS NET	Famine Early Warning Systems Network
FFP	Food for Peace
FIPS	Federal Information Processing Standards
FNID	FEWS NET Identifiers
GIEWS	Global Information and Early Warning System
GIS	Geographic Information System
HDX	Humanitarian Data Exchange
HQ	Headquarters
Hub	Learning and Data Hub
IDIQ	Indefinite Delivery, Indefinite Quantity
IDP	Internally Displaced Persons
IPC	Integrated Phase Classification

IT	Information Technology
ISSO	Information System Security Officer
JSON	JavaScript Object Notation
LOE	Level of Effort
NIST	National Institute of Standards
ODK	Open Data Kit
OLAP	Online Analytical Processing
P3	Pillar 3
PaaS	Platform as a Service
POA&M	Plan of Action & Milestones
PTA	Privacy Threshold Assessment
RAR	Risk Assessment Report
REST	Representational State Transfer
RDS	Relational Database Service
SA&A	Security Assessment and Authorization
SAP	Security Assessment Plan
SAR	Security Assessment Report
SOW	Scope of Work
TO	Task Order
UI	User Interface
UN	United Nations
USAID	United States Agency for International Development
USG	United States Government
USGS	United States Geological Survey
UX	User Experience
WFP	World Food Programme
Y1	Year One
Y2	Year Two

INTRODUCTION

The Famine Early Warning Systems Network (FEWS NET) Learning and Data Hub (Hub) Task Order (TO) 1, under the FEWS NET Pillar 2 Indefinite Delivery, Indefinite Quantity (IDIQ) contract implemented by Kimetrica LLC,[1] has the specific objective to manage, share, and facilitate the application of FEWS NET data, information, and knowledge to help achieve FEWS NET's mission to sustainably prevent food insecurity and famine. Following the project restructure in 2019, United States Agency for International Development (USAID) issued the Hub a TO to capture, document, and share the wealth of FEWS NET analytical products, datasets, methodologies and learning through three technical workstreams:

- The Data Management Platform (DMP);
- The Website Platform; and
- The Knowledge and Learning Workstream.

In Year One (Y1), the Hub split the DMP into two sub-workstreams: the Data Management System (System) and Data Platform Content (Content). The System is the application and infrastructure - including the FEWS NET Data Warehouse (FDW) and the FEWS NET Data Explorer (FDE)[2] - while the Content is the data domains themselves, as well as visualization and analytical tools, and associated system guidance and data documentation. Taken together, the Content and System are the DMP.

After obtaining operational control of the DMP during Y1 of the TO 1, the Hub conducted a series of analytical activities to inform further development of the DMP activities. This document provides supplementary information, including methodology and findings for the Hub's seven technical reviews, four prototype analyses, and six stakeholder consultations conducted during Y1. The reviews, which were authored by the Hub and external partners, include the following technical activities and associated reports, listed in Annex 1:

- User Interface (UI) / User Experience (UX) Design and WorkFlow Review of FDE
  - Annex 2. FDE Usability Recommendations Report
  - Annex 3. FDE User Interviews Report
- System Architecture Review (Annex 4)
- Gender Data Review and Gap Analysis (Annex 5)
- Compliance Review
  - Annex 6. FEWS NET Security Assessment
- Data Platform Backlog Review (Annex 7)
- Strategy Integration, including the Population Data Strategy[3] and Data Portal Strategy[4]
- Data Platform Prototype Analysis

[1] Task Order No. 7200AA19F0006 awarded effective June 24, 2019 under the FEWS NET Pillar 2 Learning and Data Hub, IDIQ Contract No. 7200AA19D00005

[2] The FDW is the backend system that contains all the raw data provided to FEWS NET implementing partners. The FDE is the frontend user interface that allows users to explore, analyze, and visualize FEWS NET data.

[3] This report is available on USAID's Development Experience Clearinghouse (DEC).

[4] This report is available on USAID's DEC.

- Stakeholder Consultations (Annex 8)

The last technical review, Stakeholder Consultations, reflects the Hub's seven interviews with stakeholders across the FEWS NET Project to understand their experience with the Data Platform and what changes they would recommend for greater usability. The Hub interviewed representatives from:

- FEWS NET Early Warning (EW) team; and
- USAID Bureau for Humanitarian Affairs (BHA) and Bureau for Resilience and Food Security (BRFS).

The above analytical activities served to inform, support, and provide nuance for two of the Hub's Year Two (Y2) deliverables submitted to USAID in July and August 2020, respectively: i) the Hub's *Annual Performance Report*; and ii) *Annual Plan for Revisions to Data Management Platform Volume 1. Core Vision and Work Plan*.

OVERVIEW OF ANALYTICAL ACTIVITIES

In Y1, the Hub engaged in a variety of analytical activities to further the goals of improving data access across the FEWS NET Project. The Hub specifically reviewed seven areas of the DMP, which are introduced in the sections below and detailed in reports within the Annexes. At the end of Y1, the Hub was in the process of reviewing the DMP's *Data Domains, Quality, Visualizations, Management and Analysis, and Ownership and Use Conditions,* and *Data Inventory* under USAID's guidance.

UI/UX DESIGN AND WORK FLOW REVIEW OF FDE

The Hub worked with external partner Development Seed, a data visualization and mapping software company, to conduct a design analysis of the FDE's UI and UX. The two-fold analysis included:

- A review of UI points of FDE pages and market research to compare FDE with other web-based data applications in both the humanitarian field and commercial market, from the *FDE Usability Recommendations Report,* available in Annex 2 of this report; and
- Stakeholder interviews related to software development, data upload and maintenance, training, data needs, and user workflows, titled *FDE User Interviews Report,* available in Annex 3 of this report.

After conducting a heuristics analysis, Development Seed found that the FDE's Usability Score was "Moderate," meaning that users should be able to use the system and complete most important tasks, but the UX could be significantly improved. Recommendations to increase this score to 100 included:

- Redesigned mockups of FDE pages to demonstrate a uniform UI that will serve as a baseline for consistent UX;
- Harmonize the many specific terms and interaction points that confused users; and
- Improve the page load speed, grouped data set presentation, and data readability.

DATA PLATFORM ARCHITECTURE REVIEW

The TO 1 requires the Hub to undertake an annual review of the overall architecture of the Data Platform, including the individual software components and the infrastructure that it runs on, to ensure that the DMP continues to meet evolving user needs. This review, titled *Data Platform Architecture Review*, is available in Annex 4 and focused on the following major functions:

- An evaluation the performance of the platform's existing technologies given its purpose, with topics including asynchronous queue management, data ingestion, data extraction, and analytic tools;
- Determining how new data management technologies may improve functionality, performance, or resource utilization; and
- Which of these components warranted further investigation or proof of concept development.

Based on the comprehensive review of all DMP systems, the Hub issued the following recommendations to improve the platform, reduce technical debt,[5] and save USAID resources:

- Upgrade the core software components, specifically Django, Postgresql, and PostGIS to ensure system longevity;
- Entirely remove and replace those ancillary libraries that are obsolete[6];
- Monitor and create additional indexes for tuning the Representational State Transfer (REST) Application Programming Interface (API) into the database;
- Investigate and implement alternatives to the existing Django Import Export code for data ingestion;
- Explore cloud-deployed ad hoc query analytical tools and apply USAID policy for selecting, implementing and supporting such tools.

DATA DOMAINS REVIEW

The current FDW comprises 10 data domains,[7] all of which were developed in the previous phase of the project. During Y1, USAID requested that the Hub complete a comprehensive review of data structures, quality, and completeness for the ten existing domains.

Throughout Y1, the Hub's monthly Sprint process prioritized these reviews during the monthly Sprint process, with any requested changes typically implemented in the same or subsequent month; for example, the Hub implemented changes to both the Market Price and Trade domains following a request from the EW team.

[5] Technical debt is the additional maintenance or development costs that will be incurred in the future as a consequence of choosing an easier approach now instead of incurring additional costs building a more long-term solution.
[6] Note this recommendation must follow USAID approval of a FEWS NET revised brand identity to be applied to all project online applications and software systems.
[7] FDW/FDE Data Domains: Price, Trade, Crop, Nutrition, Responses, IPC FIC, IPC FIPE, Population, Spatial, and Other (including Exchange rates, Semi-Structured Data, Calculated Price index, External Price Index, and Calculated Price Ratio).

At present, the Hub has no changes to domains in the backlog for Y2, although the Hub is prepared to review the Nutrition and the Response domains, should USAID or the EW team request it through the Sprint planning process.

GENDER DATA REVIEW AND GAP ANALYSIS

In Y1, the Hub conducted a review of the gender data in the DMP, titled *Gender Data Review and Gap Analysis*, a report available in Annex 5. Because the EW team is responsible for the collection and analysis of socioeconomic data, while the Hub ensures the content is accessible on the project's data and web platforms, the report reflects input from both teams.

The analysis indicated that FEWS NET has not systemically incorporated gender considerations into its analytical frameworks, modeling, or databases. The exception to this finding is the EW team's use of sex-disaggregated data in nutrition (i.e., pregnant and lactating women) and livelihoods (i.e., distinguish between male- and female-headed households). However, beyond these demographic datapoints, none of the FDW domains have sex-disaggrated indicators. A search of the FEWS NET website failed to display any documents related to the terms gender, women, woman, or female.[8]

Based on this initial overview, the Hub proposed a collaborative effort across the project to explore opportunities for providing reliable, useful, sex-dissggregated data and tools to food security analysts inside and beyond the project. Possible next steps include a more thorough analysis of gender-related data and analyses in use across the proejct, and exploring existing tools and methods for handling sex-disaggregated data, among others.

COMPLIANCE REVIEW

The Hub TO contract requires that the DMP complies with all USAID and United States Government (USG) security, internet technology, and data-related requirements, which include but are not limited to privacy protections, records management, and USAID Chief Innovation Office (CIO) system registration and audit. In Y1, the Hub contracted Insight Systems, a third party information security services firm familiar with USAID standards, to help prepare the Hub's CIO audit registration paperwork.

Throughout Y1, the Hub engaged in several analytical activities to support completion of an FIPS-199 System Registration,[9] representing the first step in USAID's Security Assessment and Authorization (SA&A) assurance process. The SA&A process will document, inspect, and evaluate the design of the DMP and FEWS NET website so they can be trusted to sufficiently secure and protect the information contained within the systems. The end result of this SA&A process is an Authority to Operate that represents the USAID CIO's formal approval to continue system operations.

[8] Note that these search terms may reside inside PDF reports, but the website's current search engine only scans report titles, not content of the reports themselves.

[9] FIPS-199 refers to the Federal Information Processing Standard's 199 publication, which are the Standards for Security Categorization of Federal Information and Information Systems.

The preparation activities the Hub undertook in Y1 are described in detail in the *FEWS NET Security Assessment* report available in Annex 6. Broadly, activities included a formal definition of the DMP's security boundary, which acts as the official description of the system that must be secured, and a readiness review to determine the gaps in meeting security controls and documentation requirements.

BACKLOG REVIEW

As part of the Agile development process, the Hub worked with USAID to review and prioritize items in the backlog—a list of outstanding items, including discovery work, tasks, improvements, and fixes—throughout Y1.[10] The resulting report, the *Hub Backlog Review*, is available in Annex 7, and is updated as of June 30, 2020.

At the start of Y1, the Hub created the DMP backlog by first reopening tasks and features identified during the previous phase of the FEWS NET project that were not suitable for implementation at the time, but which are now achievable. The Hub then added tasks to the backlog that USAID requested or the Hub identified as necessary to facilitate DMP transition. The Hub found that only two previously identified tasks should be considered for implementation during Y2 as part of the sprint process: Crop Calendars (DATA-209) and Comprehensive Metadata Documentation (RDM-10216).

STRATEGY INTEGRATION

In Y1, the Hub developed strategies to i) update the current Data Portal on the fews.net website and integrate it with the Data Platform; and ii) leverage the current Population Explorer capabilities as it relates to population data and geographic unit data for the FDE. The Hub submitted these documents, titled *Data Portal Strategy* and *Population Data Strategy,* respectively, to USAID in Y1 and expect implementation in Y2.[11]

Strategy methodologies. The *Data Portal Strategy* included discussions of the Hub's internal workflow across the DMP and Website workstreams to ensure that the information and data being displayed on the fews.net Data Portal is accurate and seamlessly displayed. In the *Population Explorer Strategy,* the Hub leveraged Population Explorer to identify processes that enable automated calculation of per capita estimates of key variables contained in the FDW, thereby enhancing FEWS NET's current and future ability to contain population and demographic data.

Data Portal Strategy results. The Hub found that the *Data Portal Strategy* is a primary catalyst to facilitate fews.net as a world-class, public-facing website. However, implementation depends heavily on the website moving from the Drupal 7 to the Drupal 8 platform, as well as consistent branding guidelines that standardize the aesthetics of the FEWS NET project.

[10] The Hub reviews backlog items periodically to ensure that they are still relevant, well-defined, and accurate given how the needs of the project can change and new priorities can emerge. Subsequently, the backlog evolves as tasks are added, prioritized, and completed through the monthly Sprint process.
[11] The Data Portal Strategy and Population Data Strategy are available on the DEC.

Population Data Strategy results. The *Population Data Strategy* is more immediately actionable, since the Hub began initial implementation during the May 2020 Sprint.[12] Given that the goal of the strategy is to ultimately pull much of the same population data that is held within the Population Explorer (www.populationexplorer.com) application, which FEWS NET relies on, the project should come to a determination as whether it is most effective to keep the Population Explorer licenses it has or rely on the Data Platform's ability to produce this data..

PROTOTYPE ANALYSES

Prototypes are preliminary models of an idea, developed as a proof of concept before the Hub commits to full implementation; it is the first step in the process of turning an idea into a feature of the Data Platform. The process to develop a prototype involves identifying ideas from the backlog, conducting research, and then demonstrating and collecting feedback from stakeholders. Thereafter, the Hub either expands the prototype into a full activity or returns it to the backlog.

In Y1, the DMP team developed several prototypes, including automation of Integrated Phase Classification (IPC) maps and automated ingestion of market price data from the Food and Agriculture Organization (FAO) API. The results of this review are in Table 1 below, which describe what work was undertaken and reviewed in Year 1, as well as what is left to uncover in Y2 for each prototype.

[12] In this Sprint, the Hub Data team loaded the LandScan demographic data to supplement the LandScan population estimates already available within FDW.

Table 1. Data Platform Prototype Review

PROTOTYPE	COMPLETED Y1 WORK	ANTICIPATED Y2 WORK
Comprehensive Knowledge Archive Network (CKAN)	The Hub reviewed the requirements for storing semi-structured data and concluded that the best approach was to extend the FDW to store the data rather than implement CKAN.	Work to implement that strategy is expected to be finished before the end of Y2.
Kobo Toolbox	In Y1, the Hub reviewed and demonstrated primary data collection platforms, including Open Data Kit (ODK) Aggregate, ODK Central and Kobo Toolbox.	The Y2 work plan will include a production implementation of Kobo Toolbox and integration with FDW.
Data ingestion via APIs	In Y1, the Hub developed the ability to define the Market Products one would want to collect and pull the price data associated with it automatically from the FAO's website into an FDW module via an API.	The team will build off the work done in Y1 and replicate the API development to connect with other USAID partners' data. Specific prospective partners are actively under discussion.
Data Visualization (Dash, Superset)	The Hub developed a very robust visualization prototype with its Data Inventory Dashboard, which was developed using Plotly Dash. This interactive dashboard visualizes the 14.5M+ data points that are held within the Data Platform,using time series, pie charts and interactive maps to display the data.	In Y2, the Hub will work with USAID to expand the visualization capabilities as well as integrate the capabilities into the Data Portal.
Automated Online Price Bulletin Report	Not initiated at this time; requires close collaboration with the EW and Hub's Web teams.	

OVERVIEW OF THE Y1 STAKEHOLDER CONSULTATIONS

In Y1, the Hub organized stakeholder consultations with representatives from USAID and the EW team to understand how FEWS NET's key consumers approach and use the data within the FDW. Discussions covered a range of topics, including how users interact with FEWS NET data on a regular basis, as well as suggested future improvements for the FDW. The full report is available in Annex 8.

EARLY WARNING TEAM CONSULTATIONS

The Hub spoke individually with three members of the EW team between April 10 and 15, 2020.

Frequently used data domains. The EW team members consistently noted that they most frequently used the Price Data domain (which includes Consumer Price Index and Exchange rate data), Crop Production data, Food Assistance Outlook Brief data, and Population data.

Use of DMP data and outputs created using the FDW. Overwhelmingly, FDW data is ultimately used to support the Decision Support Groups' needs as it relates to integrated food security analysis, with delivery through mapping data and mixed surveys, Price data, and IPC data. The EW uses the FDW to create a number of reports, by: i) running analyses on IPC and population data to show the two domains layered on one another; or ii) using prices and cross-border trade data to generate scenarios that show variability in food availability and access.

New or expanded data domains that would be of more use to the EW team. The EW team interviewees recommended the Hub expand metadata (i.e., for nutrition surveys, livelihood zones, and IPC shapefiles) and introduce new data sets (i.e., data on remittances, urban and rural income/wages, and humanitarian assistance). The EW team recommended the Hub enhance the FDE's extraction capabilities to allow customization in new formats. Lastly, the EW team proposed the Hub identify opportunities for automation. Leveraging both the FDW and new technologies could enhance collaboration between the Hub and EW teams while reducing manual work and level of effort (LOE) required for regular data tasks.

USAID CONSULTATIONS

The Hub held four interviews with staff from the BHA and BRFS staff between March 13 and 19, 2020, focusing on their current use of FEWS NET data to gain insights into their needs and experiences. Overall, USAID colleagues noted that staff across the Agency—including senior management, communications, and other analytical departments— use FEWS NET data in their everyday work.

Across these four interviews, the most used data and data domains are the IPC, Price, Livelihood and Boundary data, typically accessed through the FDE, for those with access, or via the FEWS NET Data Center available on fews.net—primarily for various reporting purposes. USAID colleagues reported frequently using FEWS NET spatial data to conduct analyses for visual presentations or aggregating data for use by senior USAID staff in congressional briefings. The Hub also found that some stakeholders prefer not to use FDE because of the level of complexity required to extract data.

Interviewees suggested several improvements for FEWS NET data gaps and delivery, including providing FDW users the ability to define the timeframe for the data that they want to extract and allowing users access to more sub-national and regional country data, as much of the data is currently only available at the national level.

FINDINGS FROM Y1 THAT INFORM Y2

Based on the results of the technical reviews and consultations described above, the Hub compiled the following findings from Y1 to inform its recommendations for Y2. The section summarizes across all the reviews above and is organized by the Hub's four primary segments of the DMP's Y2 work plan.

MANAGE, MAINTAIN, AND IMPROVE THE DATA MANAGEMENT SYSTEM

The Hub identified four primary findings related to datasets and data domains on the Data System:

User access, permissions, and data sharing policies. The Hub's *Compliance Review,* available in Annex 6, indicated that the Hub will need to update permissions and access policies as system access begins to expand beyond the current FEWS NET community. For instance, once the system is open to the public, the Hub must address several permission scenarios to ensure general users only can access the data appropriate to their permissions level. The Hub discovered priority areas to further explore in Y2.

Data extraction via REST API. Currently, the FDW delivers data extraction in a single flat or time series comma-separated values (CSV) format. As referenced in the *Stakeholder Consultations* report available in Annex 8, the Hub learned that users prefer specifying the exact types of metadata to be extracted alongside the dataset. The Hub plans to address these issues through development of semi-structured domain capabilities. The Hub recommends that USAID discuss these issues with the EW and Science teams to define the requirements accurately and collaboratively, thereafter prioritizing the issue.

User interface and experience. From a UI/UX perspective, the Hub recognizes the platform has room to improve its ease-of-use. This is particularly important for descriptive information like tooltips, which can help users understand why content may or may not display as is or the user might expect. The Hub can further address UI/UX experiences in upcoming Sprint cycles to understand specific use cases for descriptive information.

Data ingestion. From a systems perspective, the process of getting data into the Data Platform needs to be refined and agreed upon by all of the current FDW/FDE stakeholders. The process includes gathering, cleaning, appropriately tagging and uploading the data. The EW's work flow process has significant LOE implications that could be solved by use of the Data Platform. The Hub aims to grow the system and increase its use, which will require further development and documentation of the data ingestion process.

MANAGE, MAINTAIN, AND IMPROVE THE DATA PLATFORM CONTENT

DATA SETS AND DOMAINS

The Hub identified four primary findings related to datasets and data domains on the Data Platform:

Limited use of project-wide use of FDE Data. In the *Stakeholder Consultations* report available in Annex 8, FEWS NET team members noted that they frequently manipulate and sort data they wish to extract from the System. Generally, FEWS NET team members pull population domains and IPC data for reporting and decision-making support, rather than immediate analytical needs.

Improving existing domains. Results from the technical reviews and stakeholder consultations indicated that the Hub should focus on improving the most commonly used data domains currently in the System, including Crop, Livelihoods, Nutrition, Remittances, Population, Price, and Remittances. Examples of the types of improvements include, but are not limited to, updating outdated boundary data (country border shifts), adding new metadata (i.e. as new data is produced, making sure it's updated on FDW), and updating the data from a permissions/policy perspective, should there be a change from USAID or other National orgs' access requirements.

New domain opportunities. The Hub also learned the FDW has an important opportunity to store data critical to FEWS NET's early warning analyses, such as household survey data. Adding these data domains into the FDW will necessitate development of explicit data sharing agreements, clear processes for data cleaning, processing, and storage, as well as metadata requirements. The Hub recommends USAID lead a project-wide effort to identify what processes or structures currently prevent additional data domains from being added onto the system, as well as potential solutions to overcome these obstacles.

Automation opportunities. EW team members indicated that automation has the potential to reduce the LOE required to produce reporting (e.g., automated publication of shapefiles or data-centric reports, such as the Price Watch Annex).

DATA VISUALIZATIONS

The Hub identified additional automation opportunities in visualizations, particularly for price trends and commodity availability. The Hub can create visualizations through stock (i.e., template) files or can be fully automated. Automated data visualization functionalities will engage more users beyond data extraction and manipulation to analysis and exploration.

ANALYTICAL TOOLS

The Hub found that FEWS NET team members currently use tools other than the FDW to conduct regular analysis and create data visualizations. To meet these needs more effectively and reduce significantly high LOE for manual data activities, the Hub will focus on development of analytical functionalities in the FDW that improve the capabilities of team members' existing tools.

DOCUMENTATION AND GUIDANCE

The Hub recognizes that different audiences require varying levels of support to both access the system and understand the data in the FDW. The Hub will work with stakeholders across the project to develop clear data documentation and intuitive system guidance, critical for expansion of the FDW for audiences internal and external to the FEWS NET Project.

COMPLIANCE WITH USAID REGULATIONS AND INDUSTRY STANDARDS

As referenced in the *Compliance Review* report available in Annex 6, the Hub will compile the documentation necessary to register the FEWS NET DMP with the USAID CIO, ensuring it complies with USAID regulations and eventually obtaining its Authority to Operate (ATO) from USAID.

ANNEX 1. LIST OF TECHNICAL REVIEW REPORTS

ANNEX NUMBER. TECHNICAL REPORT TITLE
ANNEX 2. FDE USABILITY RECOMMENDATIONS REPORT
ANNEX 3. FDE USER INTERVIEWS REPORT
ANNEX 4. DATA PLATFORM ARCHITECTURE REVIEW
ANNEX 5. GENDER DATA REVIEW AND GAP ANALYSIS
ANNEX 6. FEWS NET SECURITY ASSESSMENT
ANNEX 7. HUB BACKLOG REVIEW
ANNEX 8. STAKEHOLDER CONSULTATIONS

ANNEX 2. FDE USABILITY RECOMMENDATIONS REPORT

The following report was commissioned by the Hub's external partner Development Seed. The report below represents an edited version of their final deliverable submitted to the Hub, presented below in its entirety. Any references to "we" refer to Development Seed.

EXECUTIVE SUMMARY

The FEWS NET Data Warehouse and Explorer provides users with access to a wide array of global food security data. Experts can search, save, upload, manage and download data series on international food markets, prices, trade, nutrition, production and other indicators. The platform also empowers users to visualize data, update data and extract data sets. This report focuses on the FEWS NET Data Explorer, which offers a "front end" for a more general audience to interact with FEWS NET data.

Overall, this evaluation finds that the user experience for the FDE can be improved with re-configuration of the application's layout, a greater focus on essential features and data, and through a number of visual changes to the interface.

TOP TAKEAWAYS

- Users need greater speed, performance and reliability for the application to be useful and preferable over direct access to locally hosted data;
- Users want to see a simplified view of the data explorer that includes the most necessary information for each data domain;
- By reviewing common patterns in top competitor platforms uncovered in market research, the Data Explorer can follow data platform best practices that will improve user experience; and
- Implementing the recommended changes from the user interface review will allow users to search for, locate, understand and download data more quickly and intuitively

METHODOLOGY

The original planned evaluation methodology included usability testing with individual end-users and data consumers. However, the users available for interview purposes were all considered "super users" of the Data Explorer platform. The majority of this audience had background knowledge of the platform's design, intentions and shortcomings, and as such did not reflect an ideal audience for usability testing. Rather, these users represented stakeholders who have worked with the software development team, data uploaders and maintainers, and implementing partners who have trained others on the platform.

Instead of general usability testing, these users answered a short set of background questions on their experience and data needs. Following these questions, they were asked to walk through their typical workflow on the Data Explorer platform while narrating their approach. This approach allowed users to describe their needs while also pointing to problem areas ripe for improvement.

In parallel to the user interviews and platform walkthroughs, we conducted market research and completed an in-depth user interface heuristic and feature analysis.

USER INTERVIEWS

We interviewed nine users from Kimetrica and Chemonics, asking them questions about their work, their standard workflow in the FEWS Net Data Explorer, and their goals and needs around the platform. These users were involved in the design or implementation of the Data Explorer platform, and as such were identified as "Super Users" for the platform. As such, these users are well versed in the data, the platform controls, and the applications of the data. Many of these users are not downloading data for analytic use, but instead use the explorer as a verification tool. Rather than conducting task analysis, the users were interviewed on their main data needs, their goals, and their past use of the platform. They were then invited to share their standard usage or workflow via screen sharing.

The users interviewed were asked some or all of the following questions:

- How do you use the FDE in your job?
- How familiar are you with the current FDE features?
- How efficiently can you accomplish your essential tasks in Data Explorer?
- How have your needs evolved since your initial use of Data Explorer?
- How easy to use are the platform controls?
- How did you initially learn to use the Data Explorer? How long did it take?
- What other tools do you consult for food security data? Which platforms are the best market comparisons and competitors?
- Do you find data in the explorer discoverable, consistent, and easy to use?
- What is your current Data Explorer workflow? And what gaps exist?
- What is challenging or blocking in the current Data Explorer platform?
- What do you do when you experience issues with the platform?
- What would make your job easier?

While the findings below detail common needs and issues identified by the super users identified, it is highly recommended that the Data Explorer team survey and speak with current and potential end-users to better understand the needs and goals that should shape the Data Explorer experience. The needs for domains include context-specific data, results, and visualizations that should be determined with input from the expert analysts using the data once the platform has been introduced to this audience.

FINDINGS

The main needs identified by the users interviewed include the following:

- I need to verify that data entered in the data warehouse is present and extractable;
- I need to see and understand gaps that exist in the data prior to downloading;
- I need to download country-specific data for analysis; and
- I need to compare data from different domains to understand trends and causation.

User flows. The standard flow identified by all users is: Search / Filter → Review → Download → Manipulate (in Excel). Users also identified the following specifications to the simplified user flow:

- The "Country" level is typically the first selection for data search and filters;
- Data review involves looking for visual gaps in the time series data;
- Data review may also include comparisons of text/numeric content between local files and the results returned by a Data Explorer search;
- Users download data in different ways; some rely on the flat .csv file, while others use the "refreshable link" in a connected excel workbook, or for more advanced purposes; and
- Manipulation of data in excel often includes recreating charts to fit the FEWS standard format.

Issues. Common issues surfaced by user interviews include the following:

- ***Page Speed*** was consistently referred to as an issue with the Data Explorer, with some users suggesting that the delays in retrieving data made interacting with internally-produced Excel spreadsheets preferable to the web platform; and
- ***Grouped datasets.*** Users spoke to the need to view all data series for a particular country, or for a specific food product. Currently, these data series are filtered first by domain, forcing users to download and compile data series from different domains for a country manually.

Data readability. Many users also spoke about the complexity of the filters and table data categories. Users spoke to the need to get data quickly and simply from the platform, but their comments suggest that the current interface provides more information than is necessary at first view.

User-proposed enhancements included:

- Enable notifications for changes in a "watched" data series that send alerts to a user's inbox;
- Allow users the ability to group linked datasets and compare data across domains manually;
- Separate or distinguish archived data series from actively maintained series;
- Provide data visualizations in FEWS NET-standard format (colors, chart type, axes, etc.); and
- Enhance its spatial capabilities, with satellite imagery overlays and map navigation.

MARKET RESEARCH

To better understand the context of user needs for the FDE, Development Seed conducted market research to define standards and best practices for the sector, and identify market gaps. The majority of platforms explored were cited by users as third party data sources or comparable platforms, as illustrated in Table 1.

Table 2. Market Research on Features Available on Similar Organizations' Data Platforms

ORGANIZATION	PLATFORM	FEATURES AND ADVANTAGES
FAO	FAOSTAT	• Home page orients users to multiple entry points to data, and provides context for the application; • Data Categories (countries, Items, elements, etc) allow for different levels of grouping data; • Sidebar contains overview and metadata on each domain, along with bulk download by region, last updated, related documents; and • Downloadable charts.
FAO	GIEWS FPMA	• Homepage dashboard visualization provides immediate overview; • Data table structure is logical, includes sortable columns; • Searchable map navigation; and • Multi-lingual support.
Agricultural Management Information System (AMIS)	Statistics Outlook	• Navigable maps, with global aggregate; and • Chart options: download in excel, print, image (multiple file types).
AMIS	Market Database	• Intuitive UI with many hierarchical options; and • Notes and definitions for each chart.
World Bank	Data Bank	• "Empty State" message indicating that users should select options; • Metadata info modal for data series/indicators (pops up on click of indicator name); • Data Display tabs allow users to change context: Table, Chart, Map; and • Multi-dimensional tabs change depending on data display context.
World Food Program (WFP)	VAM	• Categorical and hierarchically structured domains; • Integrated satellite imagery baselayers; • Navigable maps with regional drill-down and visualizations; and • Robust help center, logical visual design.
Knoema	Knoema	• Dynamic, embeddable tables and fact sheet pages; and • Multiple visualization options with data-sensitive maps.
Office for the Coordination of Humanitarian Affairs	HumData	• List page uses labels and icons that tag each dataset with notes on source, usage, downloads, and contents; • Dataset pages have "last updated" tags; • Beta/light version available for low-bandwidth settings; and • "Add data" button/option for user participation.

BEST PRACTICES

The following practices were identified as features that enable fast, intuitive searching and interaction with the data platforms researched.

1. A homepage with application context and links to supporting resources;
2. Dashboard presentations with overview information, high-level view of data;
3. The use of icons for data and navigation to provide quick reference and reduce cognitive load;
4. Navigable maps for country and/or region selection;
5. Multiple form control options for data filtering, selection, and results view (e.g., change periodicity);
6. Categorical grouping of data filters and results (Visual grouping or semantic grouping);
7. Country-dependant/country-first data selection and filtering;
8. Sortable table columns;
9. Semantic HTML for application pages, data tables; and
10. Tooltip text, definition list or glossary.

Using the list of organizations from Table 1, the matrix in Table 2 below identifies whether the organization uses each of the 10 best practice features or attributes identified in the list above. For each feature, a white dot indicated the platform used it, while a black feature indicated the platform was not using the best practice.

Table 3. Market Research Comparison of Best Practice Features

ORGANIZATION	1	2	3	4	5	6	7	8	9	10
FAOSTAT	●	●	●	○	●	●	●	○	●	●
GIEWS FPMA	○	●	●	●	●	●	○	●	○	○
AMIS SO	●	○	○	●	●	●	○	○	○	●
AMIS MD	○	○	●	○	○	○	○	●	●	●
WB Data Bank	●	○	●	○	●	●	○	●	●	●
WFP VAM	●	●	●	●	●	○	●	○	○	●
Knoema	●	●	○	●	●	○	○	○	●	●
HumData	●	○	●	●	○	●	○	○	●	●
FEWS DE	○	○	○	○	○	○	○	○	○	○

USER INTERFACE ANALYSIS

HEURISTIC ANALYSIS

A heuristic analysis provides a framework for website usability according to a common set of usability metrics. Many web platform heuristic analysis are based on Jakob Nielsen's 10 general principles for interaction design. These provide general rules, and are not specific to any particular web platform. However, they provide a helpful starting place for understanding web usability.

We conducted a heuristic analysis on the FDE to evaluate usability and uncover areas for improvement. As shown in Table 3, the FDE received an overall usability score of 59 of 100, determined as "Moderate Usability." The categories analyzed included: features and functionality, home/starting page, navigation, search, control and feedback, forms, errors, content and text, help, and performance. Note that in Table 3, the italicized text represents any relevant comments related to the rating. The scores in Table 3 below reflect the following rating levels:

- ***VERY POOR.*** Users are likely to experience very significant difficulties using this site or system and might not be able to complete a significant number of important tasks;
- ***POOR.*** Users are likely to experience some difficulties using this site or system and might not be able to complete some important tasks;
- ***MODERATE.*** Users should be able to use this site or system and complete most important tasks, however the user experience could be significantly improved;
- ***GOOD.*** Users should be able to use this site or system with relative ease and should be able to complete the vast majority of important tasks; or
- ***EXCELLENT.*** This site or system provides an excellent user experience for users. Users should be able to complete all important tasks on the site or system.

Table 4. Results of the FDE Heuristic Analysis

USABILITY GUIDELINE, BY TOPIC	SCORE
FEATURES AND FUNCTIONALITY	
Features and functionality meet common user goals and objectives.	**GOOD**
Features and functionality support users desired workflows. *Workflows are inefficient due to a complicated interface and lack of visible navigation.*	**MODERATE**
Frequently-used tasks are readily available (e.g., easily accessible from the homepage; well supported; short cuts are available). *The order of operations for user tasks can be more clearly defined, and requires a hierarchy to indicate necessary inputs.*	**POOR**
Users are adequately supported based on their level of expertise (e.g. shortcuts for expert users, help and instructions for novice users). *The platform is very complicated for novice users. Learning the platform from the application itself proves challenging.*	**POOR**

USABILITY GUIDELINE, BY TOPIC, continued	SCORE
FEATURES AND FUNCTIONALITY, continued	
Call to actions (e.g., register, add, submit) are clear, well-labelled, and appear clickable. *The main calls to action appear to be "Save New Dataset" and "Reset All," in addition to "M49/FN." Better visual hierarchy will help users navigate the platform.*	**POOR**
HOMEPAGE / STARTING PAGE	
The starting page provides a clear overview of the content, features, and functionality. *While it does not currently have a starting page, the initial view could help orient new and existing users with better navigation, and with "blank state" messages.*	**VERY POOR**
The starting page is effective in directing users to their desired information and tasks. *A complicated starting view leaves novice users unsure of how to dive into the data.*	**VERY POOR**
The starting page layout is clear, uncluttered with sufficient “white space.” *The starting view of the platform is very crowded, and lacks hierarchy.*	**VERY POOR**
NAVIGATION	
Users can easily access the site (e.g. URL is predictable and returned by search engines).	**GOOD**
The navigational scheme (i.e., the menu) is easy to find, intuitive and consistent. *The navigation menu currently resembles buttons. Domains are separate data areas and should appear as navigation to show users the search and tab features.*	**VERY POOR**
The navigation has sufficient flexibility to allow users to navigate by their desired means (e.g. searching, browse by type, browse by name, most recent). *Currently users can search, but not browse, datasets. The use case for browsing datasets may be lower priority, but should be explored further in user interviews.*	**POOR**
The site structure is clear, easily understood and addresses common user goals. *While the platform is a single purpose application, more explanatory text and detailed navigation would help new and existing users more easily accomplish tasks.*	**MODERATE**
Links are clear, descriptive, and well labelled. *Navigation links are improperly labeled as buttons. Data Management link location and styling are not appropriate to expected functionality.*	**POOR**
Browser standard functions (e.g. “back,” “forward,” “bookmark”) are supported.	**GOOD**
The current location is clearly indicated (e.g. breadcrumb, highlighted menu item). *Greater indication of current domain/search type would be helpful for user navigation.*	**MODERATE**
Users can easily get back to the homepage or a relevant start point.	**GOOD**
A clear and well structured site map or index is provided, where necessary. *While a site map is not necessary, better focus on information architecture is needed.*	**VERY POOR**

USABILITY GUIDELINE, BY TOPIC, continued	SCORE
SEARCH	
A consistent and easy to find/use search function is available throughout.	**EXCELLENT**
The search interface is appropriate to meet user goals (e.g. multi-parameter, prioritised results, filtering search results).	**GOOD**
The search facility deals well with common searches (e.g. showing most popular results), misspellings and abbreviations. *This does not apply to the search feature. However, search filters are not well hardened against error, and there is no auto complete.*	**MODERATE**
Search results are relevant, comprehensive, precise, and well displayed. *Search results are confusing due to lack of pagination or number of returned datasets.*	**VERY POOR**
CONTROL AND FEEDBACK	
Prompt and appropriate feedback is given (e.g. following a (un)successful action). *Toast notifications are adequate. More visual design could help with visibility to users.*	**GOOD**
Users can easily undo, go back and change or cancel actions; or are at least given the chance to confirm an action before commiting (e.g. before placing an order). *Save dataset form requires validation, rather than waiting for user submission.*	**GOOD**
Users can easily give feedback (e.g. via email or an online feedback / contact us form).	**GOOD**
FORMS	
Complex forms/processes are broken into readily understood steps and sections. Where a process is used a progress indicator is present with clear numbers or named stages. *No complex forms are present.*	**GOOD**
A minimal amount of information is requested and where required justification is given for asking for information (e.g. date of birth, telephone number).	**EXCELLENT**
Required and optional form fields are clearly indicated.	**N/A**
Appropriate input fields (e.g. calendar for date selection, drop down for selection) are used and required formats are indicated. *Required formats not indicated. Calendar date selection form seems only semi-operational (does not always constrain results).*	**MODERATE**
Help and instructions (e.g. examples, info. required) are provided where necessary. *"New Dataset" form lacks validation. Options for data lack detailed instructions.*	**VERY POOR**
ERRORS	
Errors are clear, easily identifiable and appear in appropriate locations (e.g. adjacent to data entry field, adjacent to form, etc.). *Lack of error messages in the data panel; nothing to indicate when/whether data is unavailable, or if search terms or filter settings are non-working.*	**MODERATE**

USABILITY GUIDELINE, BY TOPIC, continued	SCORE
ERRORS, continued	
Error messages are concise, written in easy to understand language and describe what's occurred and what action is necessary. *The toast notification error messages are well written.*	**GOOD**
Common user errors (e.g. missing fields, invalid formats, invalid selections) have been taken into consideration and where possible prevented. *Validation is needed for the "Save Dataset" form.*	**POOR**
Users are able to easily recover (i.e. not have to start again) from errors.	**GOOD**
CONTENT AND TEXT	
Content available (e.g. text, images, video) is appropriate and sufficiently relevant, and detailed to meet user goals.	**EXCELLENT**
Links to other useful and relevant content (e.g. related pages or external websites) are available and shown in context. *There are no links for datasets nor any additional supporting documentation.*	**VERY POOR**
Language, terminology, and tone is appropriate and can be understood by users. *Terms and headings are not defined. Tooltips or a glossary could help, as could more referenced help documents.*	**POOR**
Terms, language and tone used are consistent (e.g. the same term is used throughout).	**GOOD**
Text and content is legible and scannable, with good typography and visual contrast. *The application is mainly text, thus more detail to typographic contrast and hierarchy is needed to allow users to quickly scan data.*	**POOR**
HELP	
Online help is provided and suitable (e.g. is written in easy to understand language and only uses recognised terms). Where appropriate, contextual help is provided. *Help is provided, but not in context (external). Terms are not always defined.*	**MODERATE**
Online help is concise, easy to read and written in easy to understand language.	**MODERATE**
Accessing online help does not impede users (i.e. they can resume work where they left off after accessing help).	**GOOD**
Users can easily get further help (e.g. telephone or email address).	**GOOD**
PERFORMANCE	
Site performance doesn't inhibit UX (e.g. slow page downloads, long delays).	**GOOD**
Errors and reliability issues do not inhibit the user experience.	**GOOD**

USABILITY GUIDELINE, BY TOPIC, continued	SCORE
PERFORMANCE, continued	
User configurations (e.g. browsers, resolutions, computer specs) are supported. *Notes on the scope of work (SOW) indicate that the platform is not designed to be used with smaller screens, but users are still able to view the platform on any screen size.*	**POOR**
OVERALL USABILITY SCORE: 59 OUT OF 100 : MODERATE	

Given the application's specific user base and specialized functionality, the breadth of the analysis in Table 3 may appear too general for the data explorer. However, this analysis provides a useful baseline to approach the application's shortcomings. For example, in the Starting Page category, the Data Explorer ranked "Very Poor" for guideline 3 "The starting page layout is clear and uncluttered with sufficient 'white space'," speaking to the need for greater visual organization. Similarly, the guideline in the Search category "Search results are relevant, comprehensive, precise, and well displayed" was rated as "Very Poor," highlighting the need for better presentation and navigation of search results.

SPECIFIC USER INTERFACE RECOMMENDATIONS

PLATFORM STRENGTHS

The FDE collates and makes available a massive number of food security datasets. The platform allows users to search, filter and download a wide range of datasets organized into multiple domains. While the analysis below surfaces a large set of issues on the platform ripe for improvement, there are a number of usability and design strengths that can be highlighted. The application is already making use of many user-centered design patterns that provide a user experience that typically meets expectations for many of the stated user needs. These include:

- Multiple options for viewing data (table, chart, map and list views);
- Refreshable and dynamic links to both flat and time-series data;
- Ability to save data sets that a user may want to revisit frequently;
- Inclusion of powerful charting libraries, with additional information on hover;
- Attempts to standardize data reporting timelines; and
- Options to view different reporting levels.

AREAS FOR IMPROVEMENT

However, we can identify multiple segments of the platform that prevent users from easily and efficiently achieving their goals. With a large number of data series in many distinct domains, and multiple context-dependant filtering options on data dimensions, using the Data Explorer can be quite complicated. To help users quickly and efficiently find and use the data they need, the application should provide clear actions and more intuitive controls.

The specific interface regions and elements are detailed below, in relative vertical page order, with issues and suggested improvements that can be broken down into the following major categories:

- Page Layout;
- Page Actions and Interactivity; and
- Visual Design and Hierarchy.

APPLICATION-WIDE IMPROVEMENTS

As multiple users stated during interviews, the desired path to the data they need is the simplest one. Users do not want to be required to make multiple selections in order to see the filtered data. While not the case for the dataset filters, the grouping of all data dimensions in one container does make these feel monolithic and required. Users also spoke to the need to view data at a high level from the initial page view.

PAGE LAYOUT

1. Create an overview page. Much of the Data Explorer is buried "below the fold." Users want to be able to quickly understand the data available and the actions possible when viewing the Data Explorer. The present data explorer layout thrusts users immediately into the Price dataset, without providing context for the entire application. An overview page with platform-wide information and data would help orient users to the goals and functionality of the explorer. This could also take the form of a more standardized "Data Dashboard" providing high-level aggregates of the data available in the platform.

2. Allow users to view linked data series easily. Users spoke to the advantage of viewing linked datasets, like price and production data, but mentioned that it was not immediately obvious how to navigate to these combined series. These linked datasets can be highlighted visually, or shown in independant pages or sections.

For further enhancement, pivot table functionality or an in-platform data analytic tool can allow users to select, manipulate and experiment with data from within the platform, rather than needing to move between the Data Explorer and Excel for downloads and analysis.

3. Increase responsiveness. While it is true that most users will be using their desktop, it is possible that a user can open links on their phone. If application stakeholders want to restrict use to full screen devices, one solution is to display a message in place of the application when the screen is below a certain size. This solution is also available for adjusting screen size on the desktop.

Navigation, at the least, should be adjusted for responsiveness to afford more page real estate to the data explorer application, and to avoid any layout bugs when users have browser windows set to a smaller size.

INTERACTIVITY

1. Provide tooltips and definitions. Use tooltips throughout the site to provide term definitions for new users unfamiliar with abbreviations and domain-specific jargon. These can be used in filters, table headings, navigation, and help text.

2. Add descriptive text. Always provide more descriptive text when context is needed. For example, with the aggregation of admin area data and the "artificial" time series data present for certain domains, tooltips or sidebars can inform users of the special conditions surrounding this data.

3. Improve perceived speed and performance. While there is work currently underway to improve the true speed of the database search and retrieval, there are a number of improvements that can be made to enhance the perception of platform speed and thus improve the experience. Relying on more modern loading patterns, text updates and notifications that stand out, and data table pagination will all increase perceived platform speed.

4. Address loading issues. This includes the following:

- ***The loading "spinner."*** When the application is stuck loading, the rotating concentric circles are seen as an infinite spinner. To ensure "visibility of system status," a best practice in usability, the user should be informed of system status with text. Instead, have a message stating what's happening. "Loading data," or "Loading filters" etc. The platform can also use a "loading skeleton" to imply the type of data that is currently loading; and
- ***Blurred "loading" message in loading filters.*** Users may not know what this means, and the approach does not have accessibility built in. Instead, use a loading skeleton or message/UI element indicating that filters need to load, and this is in process.

Additionally, if the platform requires that all filters load before rendering the data, the page should not render each filter individually as they become available; instead, the entire page should present the loading skeleton or message so that users do not attempt to interact with the application before it is ready for use.

5. Enhance the help interface. Most users identified internal colleagues as the main source of help. The current help pages can be improved to remove this task from the busy workloads of FEWS staff and allow users to more easily diagnose issues on their own. A redesigned help interface can serve as a learning center for users new to the platform, and allow users to get quicker answers and reduce the need for support tickets, or direct contact with FEWS staff. As an enhancement, make the help interface searchable to provide users with even greater access to the resources they need.[13]

VISUAL DESIGN AND HIERARCHY

1. Establish typographic hierarchy. Ensure that all text follows a hierarchical typographic scale. The current platform does not make use of standard heading tags (<h1>, etc.) Heading tags are essential

[13] *References include:* Everything you need to know about skeleton screens; Best practices in support center in design - Salesforce; and Zendesk Help Center redesign.

for accessibility and external scraping purposes. A typographic hierarchy will help direct users' eyes as they scan the platform for the data that they need.[14]

2. Establish a color scheme. Ensure that color is consistently used throughout the application. Use color to help enhance page actions, but only when necessary. Avoid drawing attention disproportionately to buttons and actions that are not top priorities for users.

ACCESSABILITY

1. Use semantic HTML to provide accessibility and more maintainable code;
2. Use Aria roles, especially for elements that are not semantic HTML; and
3. Use field labels, and avoid using placeholders as labels.

NAVIGATION

Main navigation. The current main navigation for the Data Explorer presents distinct domains of datasets. As the main organizational filter for the FDE's vast array of datasets, the primary navigation should provide users with an easily understood and visible guide to the categories available and the user's current location. While the majority of users are likely to rely on the data explorer's search and filter functionality, a strong navigation menu is needed to organize and frame the user experience.

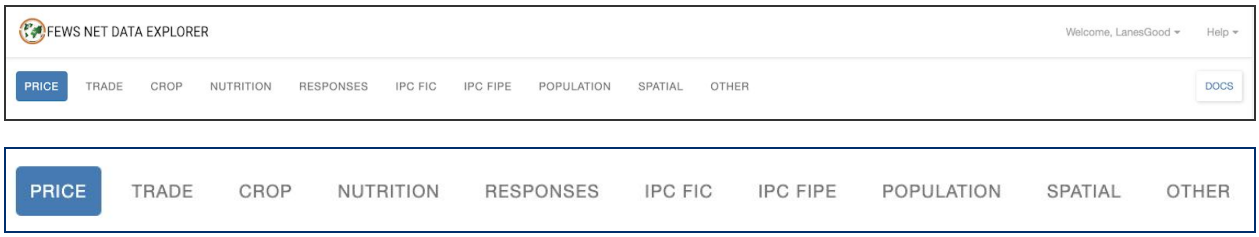

Menu items are currently styled as buttons, or navigation "pills." Visually, the combination of the filled background, rounded edges, and uppercase lettering leads users to see these menu items as interactive buttons with subsequent actions. With this unexpected pattern of menu design, it may be unclear to users that the button choice represents application navigation and affects the data in the page below. It is also unclear whether the entire application data source is changing, or if only some of the data dimensions are affected, given the persistence of certain filters. The lack of separation from the rest of the page content, and the lack of unique visual weight for the main navigation may lead the user to quickly ignore this interface. For many users, their data needs are often limited to one domain, and they may be very familiar with the area of the explorer where they can find the needed dataset. However, for novice users and for more casual browsers, a main navigation can help break up the many datasets into more digestible segments.

The current main navigation presents a confusing overview of the datasets available. The top-level domains represent a disconnected array of datasets. With many different data types, the menu should

[14] *References include:* The web is 95% Typography and People Don't Read, They Scan.

help users distinguish between the categories. New users may also be unfamiliar with the context of the data for the domains. Without any contextual text or clues present on the menu or on the domain pages, these users are left in the dark as to the meaning of menu terms or scope of the data presented.

LAYOUT

- To provide a better visual reference that navigation item selection affects the entire application, the navigation bar should be positioned at the absolute top of the page, ideally incorporating the FEWS NET logo;
- The navigation bar should be affixed to the top of the page, rather than scrolling. This is a standard navigation pattern expected by users, and will return more visual space back to the rest of the platform to use more page space; and
- The navigation may be positioned as a vertical sidebar on the left or right sides of the page.

INTERACTIVITY

- Include tooltips or glossary of terms for first-time or infrequent users. Tooltips with definition terms available on hover can help users better understand the data domains.

VISUAL DESIGN

- Rather than navigation “pills,” style navigation items as tabs that are easy to differentiate;
- Associate representative icons with data domains to help users scan a list of menu items; and
- Change navigation item capitalization to sentence case to make these titles more readable.[15]

SECONDARY NAVIGATION

Secondary navigation is lacking, despite the presence of sub-page items. Within each data domain, the explorer presents the following page functions: Search, Saved Datasets, Available & Selected data series, and finally the View Dataset Options. While a menu may not be necessary for the internal page actions presented here, more visual distinction throughout the page and top-level content hierarchy would provide a visual aid to guide users through these interactive areas. Dropdowns with in-page navigation may help users navigate to secondary actions and sections.

LINKS TO BACKEND DATA MANAGEMENT/DOC

The top level buttons “Data Management” and "Docs" take users away from the FDE and into the back-end sources and file structure of the Explorer. However, these buttons are currently more prominent than necessary. Their location in the primary navigation bar, with distinct visual styling, draws more attention than is likely intended. A more general audience seeking food security data is unlikely to also need quick access to these adjacent platforms.

LAYOUT

[15] *References include:* Tabs, Used Right and Menu Design: 15 UX Guides to Help Users.

- Add a set of top-level menu options that allows users to switch between the different applications (Explorer, Warehouse, Document management). This would make the nature of these buttons as links to external platforms more apparent;
- Situate these buttons in a platform selection dropdown/button to draw less attention to these secondary actions; and
- Include the platform selection dropdown in the top “auxiliary menu” (currently displaying the username and help) to keep it persistent, but out of the way.

INTERACTIVITY

- As the Data Warehouse and document management platforms are separate applications, they should open in a new tab. Set target=”_blank” on these buttons/links to ensure that the user’s explorer workflow is not interrupted. Currently, the button takes users away from the explorer.

VISUAL DESIGN

- Mimic the current auxiliary menu style to deemphasize these buttons, as they represent non-primary actions for users searching for food security data in the explorer.

DATA SEARCH AND FILTERS

Multiple elements of the data search and filter toolbar could be improved through better organization, visual design, and accessible form labels.

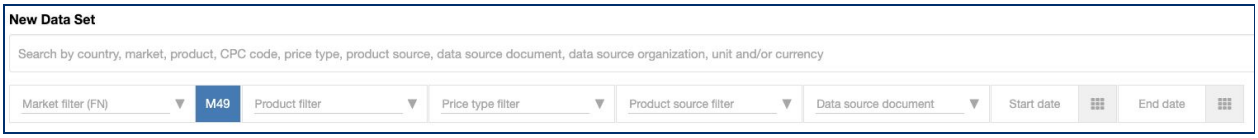

The “New Data Set” label on the search field is confusing, as most users are not creating a new data set with each search. Best practices for form input design state that placeholders should not be used to provide instructions for a user, and should not take the role of an input label. In the case of the Data Explorer’s filters, placeholders do not disappear, but when adding multiple filters the input field quickly becomes overloaded with the filter “chips” (pictured below). The filter chips are also made more prominent than necessary, with color that distracts from the rest of the platform.

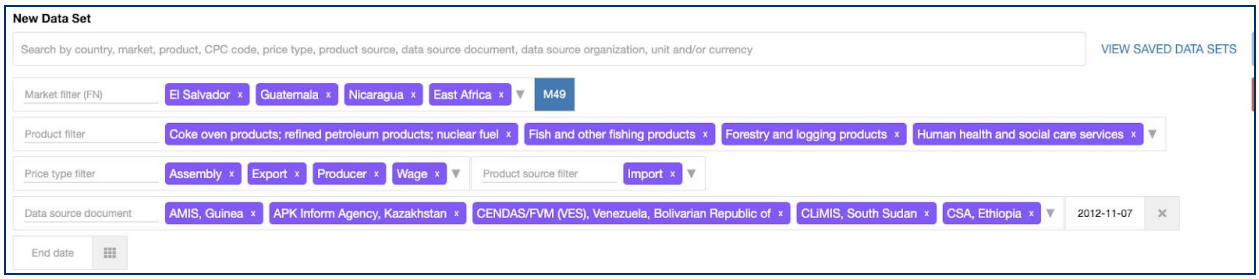

The “Market Filter” classification selection is a point of confusion to users. This button’s styling makes it stand out disproportionately to its relation to the platform, but the lack of explanatory text leads users to feel unsure of which classification they have selected.

LAYOUT

- Move filter chips below the filters, and directly above the tables. This will help users quickly see all the filters that are applied to the results table in a compiled view separate from the filter search field, and avoid the continuous repositioning of the filter inputs as more filters are applied; and
- Separate the filters: Group the filters thematically to better guide users through their application. Rather than combining filters on country, source documents, time periods, etc., filters should be organized hierarchically and within similar categories. The top-level and cross-domain filters should appear first/highest to ensure consistency between domains.

INTERACTIVITY

- Help users understand that input fields are searchable by adding a standard search icon;
- Label all filters with a true form <label>;
- Add explanatory or help text for filters, including tooltips when necessary to explain industry-specific terms for the entire filter category;
- The animated additional letter spacing for filter placeholders on focus is not an obvious pattern. Avoid changing letter spacing, and instead use a standard on-focus style like box outline or color emphasis; and
- Treat the M49/FN setting as a “utility” feature for the market filter by including this within the Market Filter’s label, dropdown, or separate filter section/widget.

VISUAL DESIGN

- Make the labels for each filter prominent, easy to understand and easy to read;
- Reduce or remove the color for filter chips;
- Remove the M49/FN button, and replace with a settings cog symbol or text link;
- Ensure that filter chips have a maximum width, and cut off text for longer filter names;
- The “X” close button for filter chips is small and indistinct. Instead of using string text “X,” the “✕” symbol ensures consistency, reducing the risk of being affected by user agent styling; and
- The checkbox icon for nested filters in the filter dropdowns should use the same checkbox icon present in the data series table for visual consistency.

ACCESSABILITY

- Adding <labels> to filters is necessary for accessibility concerns, and will help with overall page organization and visual hierarchy.[16]

[16] References include: Placeholders in Form Fields are Harmful and Don’t use the Placeholder Attribute

DATASET CONTROLS

The “Saved Dataset” actions are quite visually prominent, but represent a secondary action that should be more contained. New users may assume that they need to save a new dataset or explore past data sets to use the application. In order to help users

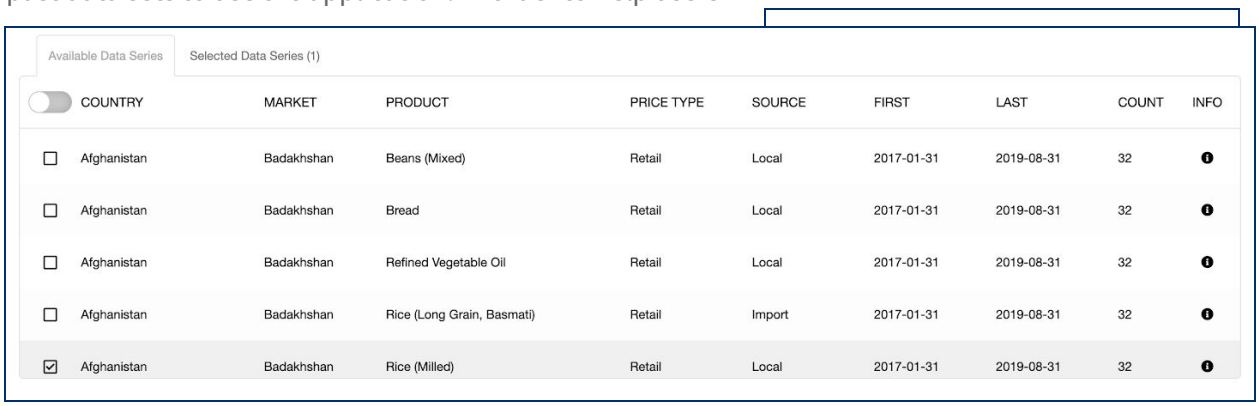

focus on the primary goal of searching and extracting datasets, this section and functionality should be repositioned and redesigned.

Not all the super users interviewed were clear on the purpose of “Saved Datasets.” Explanatory text with a description of what these are, how to create them, and when they should be used would be helpful for new users.

Saving datasets is an action for convenience and repeatable workflows. As it is not applicable for every user, surfacing saved datasets as a button on every page distracts from users’ main data searching goals. Saving a dataset is a secondary action that can be shown as an element of progressive disclosure.

LAYOUT

- Remove the current “New Data Set” label above the main search bar and rename as detailed above;
- The link for “Saved Data Sets” should be moved to a distinct primary navigation category. This allows users to access these saved datasets as a unique page within the platform, rather than including this link in the filter/search bar. Saved Datasets would thus also be platform-wide, rather than domain-specific;
- The “Save New Data Set” button should be less prominent, and should be visually secondary to the filters. As the visual flow of the platform is vertical, this button can be situated below the filters; and
- The “Reset All” button should be in a fixed location, and not be affected by filter chips.

VISUAL DESIGN

- If the Data Explorer platform remains largely devoid of color, the color of these buttons should be changed and/or muted to avoid the over-emphasis of this feature.

DATA SERIES TABLES

The data series table serves as one of the main interactive elements for users, returning the available data series that are needed for analysis and extraction. The layout, interactivity and visual display of the table can be improved to help users quickly search and process the data presented.
The data series table takes up a significant amount of space on the page. While it is useful to see all of the available data, the analysis and result tables available once data is selected is pushed to the bottom of the page and "below the fold." Designers should explore new approaches to displaying the table and results separately. Additionally, the "Selected Data Series" tab is not very evident, and yet contains the main objective for many users.

LAYOUT

- Either separate the "Selected Data Series" table into a new page section/widget, or make this tab more visually evident; and
- Separate necessary data from secondary metadata.

INTERACTIVITY

- ***Enable spatial search/navigation.*** As most users are searching first by country, spatial navigation would allow users to visually explore the available datasets. Integrate available FEWS shapefiles to permit granular navigation by administrative levels;
- ***Change "select all" UI element.*** The "select all" action is not the right UI element - a toggle should represent a global "on/off" action. The proper UI element here would be a checkbox with the appropriate help text for "Select all" to let users what action the element performs;
- ***Make columns sortable.*** Sortable table columns can help users quickly find details within an already filtered array of data series;
- ***Add icons.*** Add sort icons for sortable table columns to provide consistent visual guides;
- ***Add pagination.*** The current "infinite scroll" functionality of the data series table can be bewildering to users when browsing available data series. Pagination gives a greater sense of control - especially when setting the number of items per view;
- ***Indicate number of returned data series.*** As a best practice, indicate the number of data series present in the table. Table pagination can also indicate to a user how many items are in a table;
- ***Header text***. Consult with data analysts to ensure that table column header text is appropriate and accurately reflect viewable data (e.g., First/Last); and
- ***Metadata information.*** Ensure that the metadata available in the "Info" hover tooltip is useful and not repetitive of the data already visible in the table.

VISUAL DESIGN

- The style of text in column headers lacks contrast, but also takes away from the table text. The text should be visually distinguished from the table text styles, but also less obtrusive;
- Reduce table text size to display more data on the page;
- Reduce table row height to display more data on the page; and

- Whitespace is needed on the sides of the table to ensure that users can scroll on the page without the mouse being “trapped” by the table.

DATA RESULTS VIEW AND VISUALIZATION

As the ultimate destination for the resulting data needed by analysts, the data series results view should have high visual priority in the Data Explorer platform. Greater attention should be paid to the layout and visual design of this section to help users quickly process and understand data.

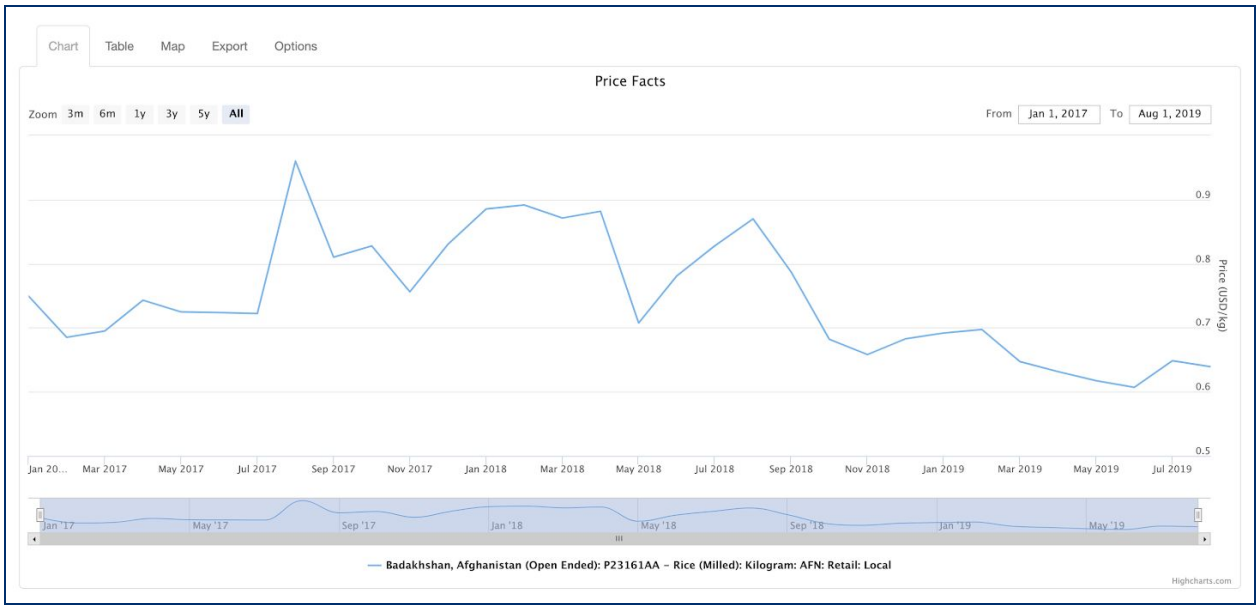

LAYOUT

- To avoid the issue of the data results showing “below the fold,” options for layouts include:
    - The results information can be displayed above the data series table
    - The result information can be displayed in a new page or new tab
    - In either a) or b), the user’s viewport position can be moved directly to the results once data is selected;
- It is unclear to new users what “Options” affects. Rather than including this on a separate tab, include as a “settings” option on the relevant tab(s);
- Make the Export/download option persistent and more visible. This action should be located on each tab, and/or be unaffected by the tabbed navigation;
- Explore separating each tab. If the table view provides the most important and persistent data, there is potential value of displaying this independently. Similarly, as the chart data is used for gap identification, explore displaying this independently in alternate layouts; and
- After consulting with data analysts, display only the necessary columns in the Table view to avoid overloading this view.

INTERACTIVITY

- Make charts downloadable in a commonly accepted file format;

- Improve map views to include useful details and visualization of more granular data; and
- Explore innovative ways to highlight gaps in the data. For example, use sparklines, provide a count of missing data points, or highlight data completeness ratings.

VISUAL DESIGN AND HIERARCHY

- If tabs are maintained, the visual style of tabs should be made more distinct;
- Explore using icons in tabs to further distinguish tabs and guide users to their function;
- Provide human-readable column titles for the Table view;
- Limit the number of decimal points to avoid overloading the data tables; and
- Follow table design best practices.

CODE IMPACT

After consultation with the Development Seed team that conducted the FDW code review, we came to the following conclusions for each feature.

Navigation. Changes to the navigation will be possible without significantly altering the page structure, code or design.

Filters. Given the complicated filters logic, the best approach for altering the filter code is to focus on the front end. Layout and visual changes may be made without affecting any interaction with the API or back end.

Tables. Many of the visual and layout changes suggested are also focused on the front-end of the application. Given the complicated workflow, any substantial changes to the flow may require a more substantial re-write to the application code. In this case, more value may be gained from rewriting the application rather than updating the existing codebase. Changes to the workflow can be done without rewriting completely, but would sacrifice some concerns of maintainability, consistency, and "react-ness" (adhering to the principles of well-designed react applications).

CONCLUSION

The FDE is already a robust web application with many modern and user-centered features built into its initial design. This UI/UX review aims to help the FEWS NET implementing organizations improve the application in preparation for dissemination to a wider audience. Once the issues uncovered by the user interviews and interface analysis are addressed with the solutions proposed above, users will be able to use the Data Explorer more quickly and efficiently to find, assess and extract the data they need for analysis. New users will be able to orient themselves more quickly to the platform's many features and data domains.

The current application provides an entry point for experienced users to search and filter for known data series and download data sets for an existing set of analytic tasks. To become more useful and intuitive for inexperienced or outside analysts, and to enable new users to perform analysis in new

ways, the Data Explorer should place greater emphasis on features that reduce the need for users to hold inside knowledge. These features include tooltips, adjustable data table parameters, contextual text and navigation. The addition of distinct page navigation, spatial navigation, categorized filters, icons and overview dashboards can all contribute to a more modern and productive application that allows users to focus on their goals, rather than on learning the Data Explorer platform.

To best prioritize the UI improvements detailed in this report, it is highly recommended that the implementing team first seek further conversations with FEWS NET data end-users and field analysts. Using the best practices detailed above and noted in the deliverables provided by Development Seed, targeting high-value structural, interactive and visual changes will improve the data search and analysis features.

ANNEX 3. FDE USER INTERVIEWS REPORT

The following report was commissioned by the Hub's external partner Development Seed. The report below represents an edited version of their final deliverable submitted to the Hub, presented below in its entirety. Note that responses to interview questions have been edited for length and consistency.

Development Seed interviewed nine users from the Early Warning and Hub teamws about their work, their standard workflow in the FDE, and their goals and needs around the platform. The interviewees were involved in the design or implementation of the Data Explorer platform, and as such were identified as "Super Users."

KEY TAKEAWAYS

The majority of users interviewed were considered FDE Super Users because of their familiarity with the platform and their involvement in designing it. These users are well versed in the data, the platform controls, and the applications of the data. Many of them are not downloading data for analytic use, but use the explorer as a verification tool. Key takeaways include:

- All users cited response speed as a crucial issue with the current platform;
- Many of the users identified the need to simplify the filters and table data categories; and
- Users discussed the need to get data quickly and simply from the platform, and suggested how to make downloads more immediate and useful.

SUPER USERS INTERVIEWED

The FDE platform Super Users include uploaders, maintainers, help desk operators, and program managers. They were not asked to rate their familiarity with the platform in great depth, as they are working on improvements to the platform with the implementing organization, or have moved from regular users of the platform to management positions.

INTERVIEW RESULTS

Note that the responses have been edited for length and consistency.

FRANK RIELY, HUB DIRECTOR, FEWS NET DATA AND LEARNING HUB

How do you use the FDE in your job?

I review all the work done on the system. As I am currently not working in an analytical role, try to take a user's perspective when looking at the explorer.

How efficiently can you accomplish your essential tasks in Data Explorer? (1-5) 3 out of 5.

It varies from data domain to domain. The product is very streamlined and efficient for the way the team expects the system to act. I have spent a lot of time producing data sets. The past emphasis has been to get data into the system. I developed features to aggregate data, and as such datasets are highly complex. The methodology of aggregations, though confident of rigor, is not transparent. Results of data download can be hard to interpret. There is a lot of work to do to simplify the data download and how the dataset is structured. The main problem is that we've often only collected the

data at the lowest admin level available. For example, we can drill down to admin 4 level and assume things add up. We sometimes only have a subset of data. We don't record this on the platform, and this leads to uncertainty in interpreting results. There is no indication to users of gaps in the data. There are no blanks for missing data; no way to know what is missing and why.

How have your needs evolved since your initial use of Data Explorer?
For many datasets, I am the only one looking at data downloads. [At this time, the FDE has only been] rolled out [to] market and trade data analysts [on the EW team]. There is demand for using population data, but it's not available in the form that analysts need. FEWS NET curates spatial datasets extensively, I don't know if they're up to date. Datasets are kept close to the chest, and then when uploaded there's no particular schedule. [The EW and Hub teams were] not in close contact [with each other at the time], [so the Hub] couldn't know when there's been a major change. There's lots of information in the system without explicit variables - Consumer Price Index (CPI), expense ratios, etc. Junk drawer of data, but the search bar can pull crop production data.

What other tools do you consult for food security data?
I often reference FAO price and production datasets, FAOStat, GIEWS. WFP is changing, it is hard to find price data - perhaps under the dataviz site. I use the DHS site for survey data, AID Data site (USAID, William and Mary), Data.gov is important, and Knoema.

HAMELMAL KIROS, HELPDESK SPECIALIST, FEWS NET DATA AND LEARNING HUB

How do you use the FDE in your job?
I am a software tester supporting the FEWS NET, testing features and serving as a user helpdesk agent.

What other tools do you consult for food security data?
I compare excel spreadsheets. Each dataset has its own template, and I use CSVpy.

What is challenging or blocking in the current Data Explorer platform?
Features on some domains are very slow. For example, in the price domain, the filter takes some minutes to load.

What do you do when you experience issues with the platform?
I contact developers to fix the issue.

What would make your job easier?
If users can get the data out of the system.

ROGER HUNWICKS, DATA MANAGER, FEWS NET DATA AND LEARNING HUB

How do you use the FDE in your job?
I am responsible for managing the technology that manages the data. I access the data often, and provide support and guidance to other users. I am most interested in learning how others can best use the data. I get involved when users report data problems like if data is not up to date, or has the wrong

numbers. My role includes trying to understand what went wrong by looking at the back end, database and data trail. I often start by looking at data in FDE. I use python or SQL queries to look at objects, and data relations.

How efficiently can you accomplish your essential tasks in Data Explorer? (1-5) 4 out of 5.
The data explorer API was designed for offline use. Not all endpoints have been performance-tuned for an interactive client. The application can seem slow and laggy. It hasn't created a comprehensive list of endpoints.

How have your needs evolved since your initial use of Data Explorer?
The data in the FEWS Data Warehouse has changed over time. Initially, we only added price data. Slowly we added more data domains. As we add each new data domain, we need to become experts on the new data, look at data in spreadsheets until we understand it. It is unlikely that there will be new datasets in the next few years; core data domains are loaded. The challenge going forward is to understand how people want to access the data. Want to understand what kind of queries and data extracts people want to get.

What other tools do you consult for food security data?
None. I have looked at Global Information and Early Warning System (GIEWS), I am aware of Humanitarian Data Exchange (HDX) - humanitarian data, and WFP data.

Do you find data in the explorer discoverable, consistent, and easy to use?
Yes, easy to discover but inconsistent; data itself is quite complicated. You need to understand what it is that you're looking at. Data is available in different time frequencies. I can look at data as a monthly average, even when it was collected weekly. Crop data is all on seasonal calendars; the nomenclature is very confusing. For example: harvest called by year planted or year harvested - refer to the same harvest, but under different names. Hard to understand.

What is your current Data Explorer workflow?
I have two ways of searching in FDE: I use filters or the search box. I prefer the search box, because it breaks down with large amounts of data.

What do you do when you experience issues with the platform?
I write a Jira ticket and assign it to someone. I hope that users send an email to the help desk. I guess they don't actually do this, but instead reach out to Nour and Linda (data people on Chemonics side). Most users probably go to them and they would create help desk call if needed. I want more direct access to users to understand gaps in functionality.

Other notes. His workflow is atypical: he enters a direct download API call in the URL bar's browser. He does data science work in Jupyter notebooks, accessing API call in a more dynamic way.

RICHARD PROVOST, DATA ANALYST, USAID

How do you use the FDE in your job?
I use FDW to upload historical crop data. I do less analysis and rarely extract data for FDE.

How efficiently can you accomplish your essential tasks in Data Explorer? (1-5) 4 out of 5. Occasionally I have speed issues because of temporary bugs. I use country filters and download buttons.

How did you initially learn to use the Data Explorer? How long did it take? I did not need to learn how to use the FDE. It seems straightforward. I have not accessed the help page. [Navigating the] FDW, help staff/help desk, [and] help information pages [can provide guidance and give you] familiarity with the site. Error messages tell you when things are wrong.

Which platforms are the best market comparisons and competitors?
EuroStat, Country specific data/annual reports.

Do you find data in the explorer discoverable, consistent, and easy to use?
I am already familiar with data, and only use explorer if upload fails.

What is challenging or blocking in the current Data Explorer platform?
Speed issues.

GARY EILERTS, FEWS NET SURGE OFFICER, USAID MANAGEMENT TEAM

How do you use the FDE in your job?
I have not done much work with the data explorer directly, but I gather datasets that go into it, and resources. I have been a manager of this for a long time. I am leaving the current position in December [to retire from USAID], and will be using Data Explorer as an end-user. These are the things that are an issue for what I want to do when I'm in the role of an FDE end-user:

- Concurrent use of satellite imagery with data sets in the warehouse. Important to preview what that will ask of the database and data warehouse. Need to ask: what do we intend to do inside the data warehouse that other users would prefer to do on their machines?
- Take data and do something interesting. Currently, I facilitate access to data by different parties. Working with NASA harvest, I know they have a strong interest in this. [There is] a huge amount of data [that is foundational to early warning analyses that should be in the DE].
- [A big UX] issue is speed of operation. The geographic filters are not cached. I did a test of country data download speed.
- Initial page display isn't what I want. I want to know basic parameters/inventory - what countries are there, and what data do we have on them. Getting in and knowing what we have is one of the first issues that users face [and it is one that needs to be improved immediately].

Concerns for future users.
The visualizations and graphs are placed at the bottom of the page of the screen; new users might not know this at all. I worry that functionality is buried. Download of data is also hidden. We need to think

about how we introduce users to the concept of associated domains. I want to see agricultural production compared to food prices.

LINDA SHEPPLER, DATABASE CONSULTANT, FEWS NET EARLY WARNING TEAM

How do you use the FDE in your job?

I have been with the project for seven years. My main work deals with markets and trade. This was the first domain that was fully functioning, and still the most active. The foundation of the data warehouse was markets and trade. I was part of the first session, to do training for East Africa. Since then, I have built up the crop production domain and continue to focus on markets and trade. I have helped with uploads for data, and I am familiar with uploads and issues. I work to upload historical data. I clean, upload, and know the different types of metadata in the DW. For Data Explorer, almost everything is on FDW. In the Data Explorer, I like to search for high-percentage increase in price, which is easy to notice and look for price data quickly. The Data Explorer extract find is easier to use - but the excel web query is easier for the Data Warehouse. I copy the export link from FDE for Excel refreshable workbooks.

How efficiently can you accomplish your essential tasks in Data Explorer?

I do not use the Data Explorer for my main tasks. Instead, I use it for searching to confirm that data exists and/or has been uploaded.

What can be improved?

I'm not crazy about all the different search fields - I just want to type a few words and get information. I am concerned about identifying prices for active data sources vs. same data sources that may have ended (change in currency, etc). I want to see price data right away after clicking on the dataset. I find the FN/M49 buttons confusing.

NOUR NOUREY, DATA MANAGEMENT ADVISOR, FEWS NET EARLY WARNING TEAM

How do you use the FDE in your job?

On top of being the data management coordinator for the FEWS NET project, I also manage the GIS component for creating maps, maintaining geodatabase and shapefiles. The data I use most on a daily basis is crop production data, from acquiring data from field staff, to extracting, processing and uploading this data through the FEWS Data Warehouse, to eventually rendering it in the Data Explorer. I manage other data domains like FIPC, Food security classification, nutrition and spatial data. FDW and FDE. I oversee price production data with Linda [Sheppler].

My interaction with FDE has been interesting. I always check the front end of FDE after uploading data. [The EW team has] not officially rolled out the product to field staff. The EW team is still working out issues with accessibility data privacy. I coordinate data completion and quality with field staff and the home office. With the team, I manage access to FDW. I [am aware that FEWS NET has] recently given access to partners at WFP and UN.

How efficiently can you accomplish your essential tasks in Data Explorer?
It depends on the data you're looking for. Price data is efficient. Not all data is efficient. [The system is] Currently improving availability and speed of crop production data. The way this data is extracted is not always as user-friendly as possible.

What other tools do you consult for food security data?
FAO GIEWS, HDX, ACLED Conflict data, United States Geological Survey (USGS), WHO Pink Sheets, WFP and others.

What is your current FDE Workflow?
I use the FDE to ensure that data is represented correctly - country names are spelled correctly, source documents are correct. I check that extracted data is correct and relevant to the domain.

What can be improved?
Filters are a high priority - many people want to click a few buttons to see the data right away, rather than go through a whole slew of filters. Filter names should be appropriate and better organized. When trying to show people a product, it needs to be well-executed so that users rely on FDE instead of excel sheet.

Cross-linkage of data between domains. For example, crop production per capita. I want to identify opportunities to link Food Insecurity and Population data. I need to start the conversation with project staff to understand what linkages are needed.

Remove redundant columns and labels on data tables, and ensure that all elements of data are relevant to the domain (e.g., date selection in crop production data). Filters, source labels, etc., should all be enhanced, and should always be "human" readable.

For the Nutrition Data Domain, fix start and end dates according to analyst's needs. Clean up "Data Source Document" field. No major visualization is possible. Multiple columns can be omitted. People still see it easier to use Microsoft Excel.

PETER THOMAS, QUALITY ASSURANCE AND METHODS ADVISOR, EARLY WARNING TEAM

How do you use the FDE in your job?
I use Data Explorer to download data sporadically. It depends on domain, (FIC). I am trying to look at trends over time, look at what has happened over time. I use the internally produced refreshable workbooks when possible.

How efficiently can you accomplish your essential tasks in Data Explorer?
For datasets with less familiarity, I am able to download and access data. For datasets with more familiarity, it is easier to directly locate local versions of datasets and manipulate from there. (All data is stored on the server locally, so can I identify changes).

How have your needs evolved since your initial use of the Data Explorer?
I am finding more use in comparing different datasets. I want to be able to compare different datasets within the same geography, for example trade and cost data.

What other tools do you consult for food security data?
WFP, FAO on the global level, and the ministry of agriculture, stats division on the country level. I search WFP and FAO if not present on FEWS. Quality of ag data + production statistics varies, so do other digging to verify data production sources. Consult source documents.

Do you find data in the explorer discoverable, consistent, and easy to use?
For domains that I use most often, generally yes. When looking to see unknowns, I find myself confused.

What is your current Data Explorer workflow? And what gaps exist?
Most is ad hoc, depending on the amount of data needed, ideally working in refreshable worksheets that can create standard charts.

What is challenging or blocking in the current Data Explorer platform?
At times, it is the speed of the platform. Even from the same place, at different times, it can take a varying amount of time for data to be available. More functionality developed on linking datasets is available, but it's not intuitive on what relationships exist and how they can be used.

VANESSA ROY, DECISION SUPPORT ADVISOR, EARLY WARNING TEAM

How do you use the FDE in your job?
I manage a team of analysts in the home office who are representatives of countries that work with national technical managers. I oversee analysts at the home office, and they are the ones who access data most often. I do pull in various data sets to assist and support analyses, or make presentations to USAID. Information kept in the Data Warehouse is the main data we're accessing regularly to do our work, though we pull data from other sources as needed.

Typically, I will see how a chart generates before downloading it to see if there are any gaps in the data series selected. If gaps exist, it’s likely not worth downloading the data selected. I do occasionally use the “Table” view of results to copy and paste data into Excel.

How familiar are you with current Data Explorer features?
It feels quite familiar. I used to be an analyst, and was accessing data for countries covered. I still use price data domain often, and also sometimes consult nutrition and crop data.

How efficiently can you accomplish your essential tasks in Data Explorer?
I can accomplish goals efficiently. I tend to pull a couple datasets interested in like retail, staple food prices, searching for info and select series. Most often I download data and will create charts I need in Excel.

What other tools do you consult for food security data?
I sometimes need to use external sources when data isn't available in the explorer. I most frequently check WFP's VAM platform.

Do you find data in the explorer discoverable, consistent, and easy to use?
Yes. But perhaps we should add a "historical data" category for unmonitored timelines. These tend to clutter the actively updated and maintained data series.

How easy to use are the platform controls?
Fairly easy, though because the platform was first designed for price data, it's rarely thought of for other domains. This may be more of an issue for communication and marketing of the platform capabilities.

What is challenging or blocking in the current Data Explorer platform?
The platform occasionally freezes. Searching by country and commodity is useful, but there is a need to weed through lots of irrelevant or unnecessary data series. Some series are redundant or no longer active.

The charts produced for the data visualization are nice, but it would be great to be able to download these charts and use them for the reports later created. Currently, charts are created manually in excel.

What would make your job easier?
Automation of charts that match the format used internally. The platform format does not match internal FEWS project standards. Monitoring tools - some sort of alert for drastic changes, either through interface or email notification.

ANNEX 4. SYSTEM ARCHITECTURE REVIEW

EXECUTIVE SUMMARY

The Hub reviewed the infrastructure and the software architectures of the Data Platform components, and considered the various key functions to identify promising new technologies for further investigation or proof of concept development.

The current infrastructure approach, using Docker containers running on Amazon Web Services (AWS) Elastic Container Service (ECS) and accessing data stored in a Multi-AZ Relational Database Service (RDS) Postgresql database, is a good fit for current and anticipated future requirements and does not require any significant changes.

To ensure the longevity of the application and avoid a build up of technical debt, the core software components should be upgraded to stay in line with the current stable versions: Django should be upgraded from 2.2 to 3.0 now, and then 3.1 and subsequently to 3.2 in Y2; Postgresql should be upgraded from 10.6 to 12.2; and PostGIS from 2.4 to 3.0. Both Django and Postgresql are currently using versions that are supported through 2022, so upgrades can be scheduled at a convenient time.

A number of ancillary libraries that are obsolete or no longer required should be removed entirely, including: Django Admin Bootstrapped, Django Admin Tools, Django Cache Machine, Dojango and Raven. All of these libraries have obvious replacements for the functionality they currently provide. However, in order to avoid wasted effort, they are dependent on development and approval of a USAID-recommended revised brand that can be applied to all FEWS NET online applications. Django Treebeard should be upgraded to use the current upstream stable version instead of a customized copy. Six[17] is a transitional package designed to help with the upgrade from Python 2.7 to Python 3.x, and is no longer required.

The existing asynchronous queue management configuration using Celery with a Redis broker is stable and functional, thus there are no compelling reasons to switch to an alternative system.

Similarly, the existing REST API approach using Django Rest Framework (DRF) for general purpose endpoints in combination with custom implementations for the performance-critical Data Point endpoints, is well-optimized and implementing alternative solutions such as FastAPI would likely only offer marginal gains.

Instead, the Hub should seek to improve the performance of the REST API by monitoring the queries executed and tuning the database by creating additional indexes where necessary, to ensure that common queries can be answered without requiring full table scans.

[17] The name, “six”, comes from the fact that 2*3 equals 6.

The Hub should also create prototypes to investigate whether columnar databases and data processing platforms can offer significant performance improvements or support similar performance across a broader range of queries. In particular, the Hub should evaluate and report on the relative suitability of AWS Redshift, AWS Athena and PySpark for acting as an online analytical processing layer to serve Data Point extracts.

In parallel, the Hub should investigate whether whole response caching is viable for common requests.
To resolve long-standing issues with data ingestion pipelines, particularly for Crop Production data, the Hub should investigate the use of AWS Glue as an Extract, Transform, Load (ETL) tool, or as a user interface for defining ETL jobs that can be run independently using PySpark. If Glue/PySpark is not suitable, then the Data Point ingestion code should be migrated to a Pandas-based solution, replacing the existing Django Import Export code.

To improve metadata management, the REST API should be opened up to allow insert and update capabilities for users with appropriate privileges.

The Hub should evaluate and recommend a cloud-deployed ad hoc query tool, subject to USAID determining a policy for selecting, implementing and supporting analytic tools. Candidates to be evaluated include Superset, Tableau and AWS QuickSight. Microsoft Power BI might be appropriate for the EW team if USAID decides that each team member can implement an independent strategy for analytic tools that does not need to be supported by the Hub.

The Hub should implement a GeoServer proof of concept to provide easier access to FDW spatial data from external Geographic Information System (GIS) tools.

INTRODUCTION

The Data Platform consists of the FDW and the FDE, together with some supporting systems that are either embedded directly within the FDW, or exposed via the FDW or FDE user interfaces. For example, the Mayan electronic document management system (EDMS) used to store supporting documents to FDW data is running as a separate system, but is accessed through FDW and FDE.

The software is primarily written in the Django web application framework, with data stored in a Postgresql database, and the FDE user interface written using the React JavaScript library.

The complete Data Platform runs on the AWS cloud infrastructure platform.

As part of the Hub's Task Order requirement to produce an Annual Revision Plan for the Data Platform, the Hub will undertake an annual review of the overall architecture of the Data Platform, including the individual software components and the infrastructure that it runs on, to ensure that they remain well-aligned with the requirements for the platform established by USAID in consultation with the various stakeholders.

APPROACH

The Hub reviewed the major components of the Data Platform from the perspective of both the infrastructure and software architectures, and considered whether existing technologies were obsolete or becoming obsolete or otherwise adding significant technical debt that needed addressing. In addition, the Hub considered the major functions of the systems (asynchronous queue management, data ingestion, data extraction, analytic tools) and considered whether new technologies were available that would be likely to offer significant improvements in functionality or performance and warranted further investigation or proof of concept development.

FINDINGS

INFRASTRUCTURE

The infrastructure is based on an ECS cluster running a variety of Docker containers, accessing data stored in a Multi-AZ Relational Database Service (RDS) Postgresql database. A small Elasticache Redis cluster is used to provide a cache for frequently requested data, and as the broker for the Celery asynchronous queue management system. The platform is deployed and managed using Ansible following the "infrastructure as code" pattern.

The current configuration is relatively modern, as it was updated when the FDW was deployed into FEWS NET's AWS account in September 2019 and takes advantage of modern ECS features such as the *awsvpc* network mode that allows individual IP addresses for each Task running on the cluster, and the *awsvpcTrunking* setting that allows much greater density of Docker containers on each ECS host.

Docker remains the best approach for deploying applications such as FDW with frequent update cycles, complex infrastructure and mixed development environments. It allows each software developer to run a full copy of the system, and it ensures that the versioned code that passed the automated test suite is packaged into an image that is deployed unchanged on the production servers.

There are alternative container orchestration services to ECS, including the market-leading Kubernetes and Docker's own Docker Swarm. AWS offers a managed Kubernetes cluster (Elastic Kubernetes Service, also known as EKS). FEWS NET is heavily invested in the wider AWS platform; it derives large benefits from Platform as a Service (PaaS) offerings such as RDS, which affords superior reliability at a lower cost than self-managed databases running in Docker containers. Consequently, Docker Swarm is not an appropriate platform for FEWS NET. EKS and ECS are both appropriate. Kubernetes is more flexible and more widely used, but also more complex. ECS is the native AWS solution and thus offers better integration with other AWS services like RDS, etc. through the use of AWS's Identity Access Management roles. Given the Hub's existing investment in ECS configuration and familiarity with the solution, EKS does not offer sufficient additional benefit to make it worth changing.

Similar arguments apply to the choice of database service. FDW depends on a Postgresql database, but in addition to the existing RDS Postgresql service used by FDW, AWS also offers Postgresql-compatible Aurora and Redshift database services. Aurora is a Postgresql-compatible proprietary database engine. It offers performance, availability and scalability advantages over RDS, for a small increase in cost. For example, an Aurora database can be up to 64TB in size, whereas an RDS one can only be 16TB. FDW is currently only about 500GB in size. Thus, the “limited” scalability of RDS is not likely to be a concern in the near or medium term.

Aurora can be run at a lower cost in a serverless environment, but with some limitations. Redshift is a Postgresql-compatible data warehouse platform that uses a columnar database structure. It is not suitable as a general purpose relational database and does not support some column types that are required by FDW, notably JavaScript Object Notation (JSON). It is generally more expensive to run than RDS and not any faster for general purpose queries at database sizes that can be handled by a single RDS node. However, it might prove useful for some specific queries as an addition to, rather than a replacement for, Postgresql on RDS. Given that RDS is not close to reaching current RDS scalability limits, and all the environments (dev, stage and production) share a single database cluster, there is no compelling reason to switch to either Aurora or Redshift. Furthermore, moving to one of the AWS proprietary databases for the primary data store would make the development environments more complex to maintain and less reliable at identifying errors caused by database interactions.

The Data Platform architecture is documented and managed as Ansible code held in a Gitlab repository. The AWS-native approach to infrastructure-as-code would be to use AWS CloudFormation to control the infrastructure. Ansible has the advantage of being able to control the internal configuration of individual servers in addition to the configuration of AWS resources at the expense of being more complex and harder to use for managing some types of AWS resources. However, the Hub team has already solved the problems related to using Ansible with the full infrastructure stack, and those components are unlikely to change in the medium term. Thus, Ansible remains the best approach for managing the infrastructure.

SOFTWARE COMPONENTS

DJANGO

FDW is currently running on Django 2.2, the current designated Long Term Support version of Django, which is supported through Q1 2022. The current stable version of Django is 3.0. Although FDW is on a currently supported version, the Hub’s experience with the upgrade of the Data Warehouse from Django 1.8 to 1.11 LTS and then to 2.2 LTS suggests that waiting for 3.2 LTS to be released in April 2021 and then upgrading directly from 2.2 to 3.2 may require significant effort. Although Django itself is relatively easy to upgrade from one LTS version to the next, FDW relies on a large number of additional packages from the wider Django ecosystem. This is typical for complex Django applications: the rich ecosystem of additional packages is the reason that Django is such a productive development environment. However, the wider ecosystem is not as well tested or as clearly

documented as the Django core libraries. Consequently, the upgrade from 2.2 LTS to 3.2 LTS will involve significant testing that will likely identify a number of packages that are incompatible with the new version and require either custom fixes or identification and implementation of a replacement library with similar functionality. For the upgrade to Django 2.2 this took considerable effort.

It is difficult to schedule a major upgrade requiring significant LOE while continuing to provide regular functionality improvements through an Agile Sprint process. FDW and FDE follow a monthly release cycle, with new features being released approximately every four weeks. This means that a system upgrade that takes longer than four weeks to complete has to chase a continually moving target in terms of required functionality, which adds to the level of effort required. For systems upgrades that take more than four weeks to complete this means that the upgrade has to chase a continually moving target in terms of required functionality, which adds to the level of effort required.

Exclusively using the Long Term Support versions of software is generally the best policy for applications that are going to be developed and implemented and then moved into "maintenance mode," where a smaller team maintains the application at a much lower level of ongoing effort. However, for applications like FDW that have a permanent software development team and a regular release cadence that is much shorter than the release schedule of the underlying components, it is frequently much better to track the stable version of those components rather than the LTS version. This approach provides earlier access to new features in the underlying component and avoids a major drop in the pace of overall software development while the LTS upgrade is performed.

Consequently, the Hub recommends that FDW is upgraded to Django 3.0 now, and then to 3.1 when it is released toward the end of 2020, and subsequently to 3.2 LTS in Q2 2021.

ADDITIONAL DJANGO PACKAGES

As described above, FDW relies on a wide selection of packages from the Django ecosystem. To ensure that the Hub can continue to upgrade to new versions of Django it needs to review those packages and remove ones that are no longer maintained. The Hub can "adopt" packages that are unmaintained but provide critical functionality to FDW, and take responsibility for upgrading them to current Django standards. The Hub currently does this for the pypi.org/project/django-binary-database-files package. However, the preferred approach is to use packages that have broad adoption, and an active community providing upgrades and improvements. Consequently, the Hub should periodically review packages that are "abandonware" and/or have otherwise become obsolete, for example, because the Django core now offers equivalent functionality, and remove them.

This review recommends removal of the following packages from FDW:

Django Admin Bootstrapped provides a Bootstrap theme that is used by FDW. Originally this theme offered a better user experience than the native Django functionality and was easier to customize. However, Django has updated its own built-in interface since, whereas Django Admin Bootstrapped has not been updated since 2015. However, when the existing theme is removed, the look and feel of

FDW (but not FDE) will change, and there will undoubtedly be some work to do to customize the default theme. Consequently, it makes most sense to perform this work once there is an approved FEWS NET brand and style guide that can be used project-wide, including on the website and FDE. This will avoid multiple user interface changes in relatively quick succession and the duplication of effort that would entail. The Hub is waiting for guidance from USAID on the process for defining a project-wide brand that can be applied to all online properties.

Django Admin Tools provides the current FDW landing page and the main navigation menu. It adds complexity for altering the menu, and the Data team does not use the landing page functionality for collapsible panels, etc. It has only infrequent updates (once or twice per year) and is not explicitly compatible with Django 3.0. Removing it would require the team to implement a new base template for the application that contains the application menu. As with *Django Admin Bootstrapped*, it makes more sense to wait for the publication of a project-wide brand before starting this task, to avoid duplication of effort.

Django Cache Machine is used to provide caching for frequently accessed query results, e.g. look ups of Country names from their ISO 3166 alpha-2 code. Cache invalidation is complex and Cache Machine has been responsible for a number of bugs that were hard to diagnose but were ultimately caused by stale cached data. Consequently, the Hub makes minimal use of the Cache Machine and could remove it and replace it only where there is a noticeable impact on performance. The Hub could probably easily replace it with Django Cache Memoize, which the team is already using for more general caching of function results.

Django Treebeard is a tree implementation package for Django that is used to manage the hierarchical models, specifically the Classified Product and Country Group hierarchies. FDW is currently using a fork of the upstream repository that was created by Steel Kiwi. Kimetrica subcontracted Steel Kiwi to assist with the migration of FDW from Django 1.8 to 2.2. That fork is very out of date compared with the upstream master branch, which now supports Django 2.2 directly. FDW contains some workarounds for Treebeard bugs related to character-based primary keys. To avoid future maintenance costs associated with updating the forked repository, the Hub should replace the current version with the most recent stable tag from the upstream repository (as published in PyPI), and then determine whether the workarounds in FDW code related to Treebeard are still required.

Dojango is used to provide the scrolling grid for visualizing Data Points that is accessible from the various FDW management screens for Data Series subtypes. This functionality was made redundant by the introduction of FDE, which provides a much better interface for searching and browsing data. The Data team should remove the Browse button on the Data Series screens, and then remove this package and the associated Dojo static files. Similarly, the team should remove the Chart button that gives access to the Highcharts chart implementation, which is also replaced by superior functionality in FDE.

Raven is the legacy client for the Sentry.io error reporting system that is used to capture errors from production servers. FDW should be updated to reflect Sentry's replacement of this package.

Six was a utility package designed to help projects manage the transition from Python 2.7 to Python 3.x. FDW now runs on Python 3.6 and this package is no longer required.

POSTGRESQL AND POSTGIS

FDW currently runs on a Postgresql 10.10 database, with PostGIS 2.4.4 installed to provide spatial functionality.

The most recent version of Postgresql available on RDS is 12.2. Postgresql 11 contained performance improvements related to parallel query execution that are likely to be beneficial to FDW. Therefore, FDW should be upgraded by creating a small second RDS database instance using Postgresql 10.10 and switching the *fdwdev* environment to use it. The database can then be upgraded to Postgresql 11.6 and comprehensive testing can occur. Once stability has been ascertained, the instance can be increased to the same size as the production database (currently db.r4.2xlarge) and benchmarking can be performed to establish any performance improvements. A similar testing cycle can then be used to test an upgrade from Postgresql 11.6 to Postgresql 12.2: reduce the database instance size, upgrade RDS to 12.2, perform stability testing, and increase the instance size to match production and perform benchmarking. See Amazon's documentation for more details about upgrading Postgresql RDS instances.[18] Postgresql 10.x is supported through November 2022, thus upgrades can be performed at a convenient time.

In parallel, the Docker containers used in development can also be upgraded to Postgresql 12. The Data team is currently using the mdillon/postgis Docker image, which runs up to Postgresql 11 and PostGIS 2.5. This is the most widely used PostGIS container available, but has not been updated in the last 12 months. The PostGIS project now provides an official Docker image, and so the team should switch to using those containers, specifically the *postgis/postgis:12-3.0* image.[19]

REDIS

FDW is currently using Redis 5.0.6 which is the most recent version available on AWS Elasticache. Redis meets all the Hub's current requirements and there is no need for any changes.

NEW TECHNOLOGIES

The Hub reviewed the key functions of the Data Platform and investigated whether more modern tools or approaches would be likely to offer cost-effective improvements to the existing capabilities.

ASYNCHRONOUS QUEUE MANAGEMENT

FDW uses the Celery distributed task queue for managing asynchronous tasks, such as sending emails, calculating statistics, pre-caching expensive queries and performing large uploads. FDW uses Redis as the task broker (the component that holds the queue of tasks waiting to be executed).

---

[18] The upgrade will probably need to use `SELECT * FROM pg_available_extension_versions;` and upgrade PostGIS prior to doing the engine upgrade, as described in PostgreSQL.ExtensionUpgrades

[19] Switching the Postgresql major version will require developers to create their test databases from a backup. See RDM-7784 for a more detailed discussion.

Celery is the de facto solution for distributed queues in Python software projects, but has a reputation as complex to set up and manage. Celery recommends using RabbitMQ as the broker in preference to Redis. In recent years a number of alternatives have emerged, including RQ and Dramatiq, that use Redis as the recommended broker.

FDW uses the Redis broker with Celery because FDW depends on Redis already as a distributed in-memory cache and thus using Redis avoids the requirement to support RabbitMQ as an additional service within an already complex architecture. It is also possible to use Django itself as the broker, but traditionally that was seen as only suitable for small deployments. Currently a Django Broker is not supported or recommended by the core Celery project.

Celery is primarily designed for situations where thousands of concurrent tasks need to be managed and where high throughput and scalability are critical. FDW uses Celery to handle long running tasks but at relatively low volumes, and so the ability to handle high concurrent task volumes is not relevant. RabbitMQ offers no benefits to FDW over Redis as a broker, and even the Django broker would probably be adequate.

RQ and Dramatiq are alternative queue systems that offer much simpler setup and configuration, and use Redis as either the only broker or the default one. They both offer periodic tasks (a key FDW requirement) using add-on packages. There are a number of other packages that could be used but which have less traction and/or functionality, including TaskTiger, Huey and Django-Q.

Celery offers more advanced functionality than any of these choices, for example workflows consisting of chains of related tasks, but FDW is not currently using any of those features. Celery also offers a Django Admin interface for managing scheduled tasks, including the frequency at which data ingestion from remote APIs runs. Although django_dramatiq would make it easier to switch from Celery to Dramatiq in existing Django-based projects, it still would not offer the level of integration provided by the more established Celery.

Dramatiq is a better and more modern alternative to Celery, and if the Hub were implementing a new project now the team would choose it because of the lower level of effort that would be required to develop and implement the system. However, FDW already has a working Celery implementation and considerable effort has been invested in making the queues robust and efficient. Although it is more complex than is ideal, the current setup is reliable and well-documented. Consequently, migrating to a new package such as Dramatiq will inevitably cost more in the short term as the team learns a new package and migrates existing code. There would be a long term benefit in reducing technical debt, particularly in reducing the learning curve for new developers trying to troubleshoot problems and who have not previously worked with complex Celery environments. However, in the current situation, with no major changes to the asynchronous queues anticipated, the cost benefit analysis is in favor of remaining with the existing setup.

EXTRACT PERFORMANCE

Application layer. The performance of the REST API used to support data extracts from FDW is a critical component in the overall performance of FDW. The original API was implemented using DRF since it was, and remains, the default package for enabling REST APIs for Django projects.

However, DRF is relatively slow compared to more modern alternatives such as Serpy and FastAPI. It is designed for easy implementation of read and write APIs for relatively low volumes of data, rather than fast extraction of multi-million row data sets such as the FDW Crop Production data.

Furthermore, the original REST API design was focused on the needs of Microsoft Excel as the primary client, and it was tailored for downloading large amounts of data for further analysis offline. The introduction of the FDE as a browsable interface for exploring FEWS NET data has led to an increasing emphasis on fast response to queries using paginated responses that return the initial rows of data quickly, and additional rows of data in subsequent requests.

Significant effort has been put into the REST API, including the evaluation of Serpy as a replacement for the DRF serializers and the subsequent development of custom serializers based on the Pandas numeric library that combine fast serialization with data aggregation.

Consequently, although it is possible that introducing FastAPI would improve the performance of the API, the gains would probably be small and significant development time would be required to recreate the data aggregation functionality offered by Pandas.

In most cases, the limiting factor on API performance now is the time required to execute the database query and return the results to the application server.

Database layer. In recent months, efforts have been made to improve the performance of the Crop Production queries in particular; the Crop Production data domain has more than five times the number of Data Points than the next largest domain. These efforts have included the introduction of the Crop Production Facts materialized view. This approach has been very successful at improving the response times for the queries that the view has been tuned for. However, an unindexed query against the crop production data domain takes approximately six minutes to execute, regardless of the number of rows returned. Therefore the Hub should monitor the queries executed against the view, and create additional indexes as necessary in order to tune common queries.

A similar approach to additional indexing can be taken for the other materialized views, including the Market Price Facts and Data Point Facts views.

The exact queries to tune could be identified by configuring the Postgresql parameters to log the SQL for any query taking longer than, say, 15 seconds, and configuring Datadog to read the Postgresql logs from RDS. DataDog could then be used to report the slow queries that are frequently requested.

However, with such a large dataset (there are more than 12 million Data Points in the Crop Production data domain) it is inevitable that the very flexible query structure allowed by FDE will result in some queries that do not use the existing indexes and are consequently slow.

Online Analytical Processing Layer. Large data warehouses have traditionally used an online analytical processing (OLAP[20]) engine to speed query response. Frequently, the OLAP engine was a second database that contained a copy of the data from the main database, restructured in such a way that common queries could be answered more efficiently. FDW currently uses materialized views to perform this function, as described above.

Increasingly, modern data science and analytic workloads are moving from relational databases such as Postgresql to distributed columnar[21] databases such as Cassandra, Druid and HBase. Amazon offers its own columnar database, Redshift, which was originally based on Postgresql, and consequently retains SQL compatibility. It is possible that replacing the materialized views with tables in a columnar database or other analytic engine would provide significant performance improvements, particularly for ad hoc queries.

Redshift is the most appropriate columnar database for FDW, because of the Hub's familiarity with Postgresql (which is used for the main FDW database) and because AWS provides Redshift as a managed service. A single node Redshift cluster using a dc2.large instance would cost $150 per month for a reserved instance with a one year term. The Django application would treat Redshift as a second database, that is only used for the specific models, such as CropProductionFacts, where a columnar approach is likely to result in significant performance gains. The Redshift tables would use interleaved sort keys to provide increased performance across a range of filters. In a development environment, the Redshift database could be approximated by continuing to use materialized views in the main database, and configuring the environment variables to use the same Postgresql host parameters for both databases. The team could also create a second database in the main Docker Compose file using an older Postgresql version, to enable testing of the scheduled tasks that would refresh the Redshift database. The fact that "real" Redshift cannot be run in a Docker container in a local environment will mean that it is no longer possible for developers to have a full working copy of the application, but the performance benefits may be worth it. Data would be copied from the corresponding Postgresql view into Redshift on a schedule, using a Celery task to execute a query on the main database to copy the data to a table in Redshift via a Postgresql Foreign Data Wrapper or dblink[22]. The Hub should manually

[20] See Online Analytical Processing
[21] See Column-oriented DBMS
[22] See AWS Data Blog. Although dblink is required for pushing queries from Postgresql through to Redshift, for copying data a Foreign Data Wrapper (typically, but confusingly, abbreviated to FDW in Postgresql documentation) may be preferable. Django would access the Redshift database directly for data extraction via the REST API.

create a small Redshift cluster and copy the candidate tables (Crop Production Facts and Market Price Facts), and then test query times for a wide range of queries identified by the database tuning described above, including those ad hoc queries where it was determined that additional database indexes were not appropriate.

A similar SQL-based approach would be to use AWS Athena to query the data after exporting it from RDS to S3 in Parquet format using a Celery Task. AWS Athena is a serverless interactive query service, and so the cost for using it would likely be minimal. Parquet is a columnar on-disk storage format that supports predicate push down with Athena. Performance is likely to depend on the identification of a suitable partition scheme for the data, and the proportion of queries that make good use of the partitions. Unlike Redshift's interleaved sort keys there is no easy way to support queries against a wide range of predicates. Athena is based on Presto, which could be used as a proxy in local development environments. Athena would be lower-cost than Redshift, but would probably offer lower performance.

A possible alternative to a columnar database for reducing query execution time would be an in memory analytic engine such as Spark. Spark can be considered as a server-based, distributed version of Pandas. As such, it isn't a database and doesn't have any built-in data storage capabilities. However, it could read the relevant data into memory on startup and then expose it as a Global Temporary View. The fact that it holds the entire data set in memory and doesn't require any disk access in order to answer queries should give it a clear performance advantage over disk-based platforms such as Athena and Redshift. Although Spark is largely written in Scala it is a widely used component in Python data science infrastructures and offers excellent support for Python via the PySpark library. A custom Django database backend could be produced that uses Spark SQL to run queries against the Global Temporary Views and then implemented in FDW as a second database, as described for Redshift above. The data could be refreshed by a scheduled Celery task that executes a Spark job to read the data from Postgresql to Spark using JDBC and then calls createOrReplaceGlobalTempView to expose the data. This approach is less obvious than Redshift, in that it requires a custom Django database backend. However, the backend would need to support a limited range of SQL commands on a read-only basis and it may be straightforward to create a suitable backend by subclassing django.db.backends.base.base.BaseDatabaseWrapper. A Spark-based approach has the advantage that it could be added as a Docker container to the existing configuration quite easily and would meet the objective of allowing a complete replica environment to run in development, while allowing us to scale to using PySpark with a standalone cluster on AWS Amazon Elastic MapReduce if necessary. This will be a particularly attractive solution if we use PySpark for ingestion pipelines (see below).

The Hub should investigate Redshift, Athena and PySpark and report on their suitability as an additional tool for improving extract performance.

Both PySpark and Redshift will probably perform well aggregating data up the Classified Product hierarchy based on the CPC v2 code and the Administrative Hierarchy using Common Geographic Units. However, they may struggle with resampling frequencies, e.g. Weekly to Quarterly.[23]

The Market Price, CPI and Exchange Rate Value data domains use SQL window functions to report additional statistics on each data point, e.g. price last year, 2 year average, etc. Those statistics are currently calculated by the database and cached using a materialized view. As we investigate both PySpark and Redshift we should consider whether to continue to calculate those values in Postgresql or whether we should just offload the calculations to the analytic engine.

The Hub has also considered using Dask to run to scale up the existing Pandas query infrastructure, running filtering in Python code rather than SQL. However, it is not obvious how we would read and refresh the data so it is always available, and Pandas/Dask-based filtering is probably harder to integrate than the SQL-based Spark.

Similarly, we have considered Data Series-level caching of data and then performing filtering against the cached data using Pandas, but as a custom-built solution this seems likely to introduce more technical debt than a Redshift or PySpark-based solution.

Finally, we have considered using an-memory Sqlite database on the application servers answering REST API requests, and while this has the benefit of supported integration with the Django ORM it is also not clear how we would refresh the data in a running server, and the system is likely to be less performant than a PySpark or Redshift one, as well as less mainstream.

Cache layer. Regardless of the solution implemented to handle ad hoc queries, we can also improve performance by identifying common queries, e.g. the extract of all Crop Production data for a single country, and pre-cache that response using a Celery Task so that it can be served immediately upon request. This solution is worth investigating alongside the evaluation of alternative platforms, especially if the solution can join cached data from multiple countries into a single response, e.g. in response to a request for data for a specific region.

INGESTION PIPELINE PROBLEMS

Currently, FDW uses the Django Import Export library to ingest data. This library was originally designed to make it easy to import (or export) data matching a Django model via Import and Export buttons in the Admin list screen for the model. It primarily operates online, and is manually triggered by uploading the file, and includes a preview screen of changes.

This approach was originally chosen for ingesting the Market Price data, and was subsequently extended for the other data domains. The main advantage of this approach was that loading using Django models ensured that all the data validation code for the models was automatically applied to

[23] It is worth examining the Superset SQL Alchemy query code to understand how they implement resampling against SQL databases.

the data as it was ingested. The later implementations differ slightly to the Market Price implementation, but all implementations suffer from the following five problems:

Memory usage. The design inherently creates a Django model instance for each row being imported. If the model contains foreign keys, then model instances for those keys also have to be instantiated. For models like Crop Production with a large number of parent models that have to be created (up to Data Source Document) this can result in very large memory usage. For example, a large Crop Production Import can take more than 80GB RAM in the server running the import.

Speed. Similarly, the approach of processing row by row and instantiating Django models to check foreign keys is very slow. We have improved speed by caching fetched parent models, but this adds complexity to the code, and does not completely resolve the issue of the large numbers of database queries required to validate models.

Limited error reporting. The library is designed for online use, and so although we have extended it to support offline usage via the Data Uploads screen, the error reporting is limited and it can be hard for users to understand how to resolve metadata errors or work out which rows in the data are causing them, because un-crosstabbing uploaded files can result in a different number of rows in the processed file to the uploaded one.

Lack of support for data cleaning. The library is designed primarily for loading clean data in a single transaction. It doesn't provide any support for data cleaning or normalization. FDW has implemented this functionality within the classes used to process imports, but it adds considerable complexity and makes it hard to debug errors.

Long upload time for large files. The row by row processing in a single transaction is particularly cumbersome for dealing with very large upload files where it can take up to 4 hours to process a file. Resolving metadata issues can be an iterative process, and the requirement to wait 4 hours after each correction means that it can take days to resolve errors and load the file.

INGESTION PIPELINE SOLUTIONS

In order to resolve these issues, the existing ingestion code should be replaced with a purpose-built ingestion pipeline that implements standard ETL approaches to break the ingestion process into separate steps for cleaning, normalizing, validating and loading data so that users can get clearer feedback without significant delays, and that validated data can be imported into FDW with reasonable memory and processing resource utilization. Having a separate ETL pipeline will require duplication and ongoing maintenance of the data validation applied by the existing Django models, but this cost is worth it given the likely reduction in resource consumption, load times and technical debt.

There are a number of different approaches that warrant further investigation.

Replacing existing code with Pandas. The simplest approach would be to replace the existing data ingestion code with a pipeline based on executing a sequence of data normalization, cleaning, validation and load steps using Pandas. It is possible that we will use Luigi to structure the pipeline to provide more easily understood code units and produce intermediate outputs such as the cleaned, normalized but invalid data. This approach has already been proposed by the Hub and is in the FDW product backlog.[24]

AWS Glue. A second approach would be to use AWS Glue as an ETL service. AWS Glue is a fully managed ETL service. It can discover data resources already available on AWS, such as the main FDW database and uploaded files stored on S3, and provides an interactive user interface for building data transformations. The resulting ETL pipeline is implemented in PySpark and can be downloaded for further customization using a normal development environment. AWS Glue is paid for by the second (with a 10 minute minimum) and so it offers potential to lower costs overall because we currently run sufficient ECS servers that one is always available to run the existing Data Upload tasks. In development environments we can investigate using a PySpark container to approximate AWS Glue. It is also possible that PySpark will be sufficiently efficient that we can run the container in ECS instead of AWS Glue, and just use Glue as a user interface for building ingestion pipelines.

Python pipeline. A third approach would be a local Python pipeline built on an ETL-specific library such as petl.readthedocs.io or bonobo-project.org rather than Pandas, but both Pandas and PySpark seem more flexible and we should investigate those approaches first.

In addition to custom ETL pipelines for dealing with bulk data, mainly Data Point and Data Series ingestions, FDW should enable write access via the existing REST API for metadata management. Although it is likely that metadata imports via the FDW user interface will continue to be required for the EW team members responsible for metadata management, the Hub will increasingly want to be able to automate the management of metadata from Jupyter notebooks and similar scripting environments. The existing DRF REST APIs are restricted to read-only through configuration, and the level of effort required to enable them for accepting POST requests that save data into the database is low. The REST API can reuse the Django model permissions that are used by the existing administration screens to determine access rights.

ANALYTIC TOOLS

FDW has always followed an “API first” approach to analytic tools, exposing the data to analytic tools using a client-agnostic REST API rather than specific extract functionality. This REST API powers the refreshable Excel workbooks used by the Field Office EW team members as well as the FDE user interface. The API is browseable and implements the OpenAPI[25] documentation standard. This means that FDW is well prepared to support a wide range of different analytical tools.

[24] See DATA-64

[25] An API documentation and interaction standard previously called Swagger

Any new technologies that we introduce to improve Data Point extract performance, such as Redshift or PySpark will still be made available via the REST API and will be client-agnostic.

This approach means that FDW will be well-positioned to support users who need to use their own tools for various reasons such as integration with other data or existing investments in skills or infrastructure. The Hub assumes that the Science team members and external academic researchers will be in this category.

The Hub will continue to provide custom visualizations to support EW team members and to add interactivity to the FEWS NET website once the Drupal 8 upgrade is implemented.

The Hub likely has opportunities to improve the tools available to EW staff by selecting and implementing advanced ad hoc query tools. The Hub is waiting for direction from USAID on the process to be followed for selecting analytic tools for use project-wide. Some obvious candidates for evaluation include:

- ***Tableau*** is the market leader and an obvious, if expensive, candidate.
- ***AWS QuickSight*** is likely to be considerably cheaper than Tableau and will have direct access to FDW data, as well as providing support for user-uploaded spreadsheets, etc.
- ***Superset*** is an open source ad hoc query and dashboard tool that Kimetrica has used on other projects. It would not incur any additional costs as it would run on the existing ECS Cluster and has no associated license costs. Superset has the additional benefit of supporting API queries directly, and thus it enforces application-layer security roles automatically. Tableau and Quicksight are designed to access data directly from the database, which may require additional work to restrict results according to user roles, or restricting access to the analytic tool to those users who are allowed to see all data.
- ***Microsoft Power BI*** is a good technical solution, and might be suitable for EW its own internal IT team can implement and support it. However, the Hub would not be able to provide support beyond the REST API layer.

Given the budget, training and support implications of project-wide analytic tool choices, the Hub will need to wait for USAID to confirm the policy on tool selection before continuing.

FDW is based on a Postgresql database with the PostGIS extension. Consequently it has good support for spatial data. It already identifies individual Features from uploaded shapefiles and can extract Geographic Units as well as Data Series and Data Points in GeoJSON format via the REST API.

However, the spatial data can be difficult to access for users more familiar with traditional GIS tools. Both the Science and EW team members would likely benefit from a browsable interface that lets them integrate FDW data with the existing spatial tools. The Hub should investigate implementing GeoServer as a front end to the existing spatial data stored in FDW to support easy integration with ArcGIS and the USGS/Climate Hazards Center GeoEngine.

There are alternative spatial platforms that could be investigated, such as GeoTrellis, but as GIS functionality within FDW is focused on providing a spatial reference for tabular data, rather than spatial data processing, the additional complexity is unlikely to be worth it.

ANNEX 5. GENDER DATA REVIEW AND GAP ANALYSIS

INTRODUCTION

The Hub's specific objective is to manage, share, and facilitate the application of FEWS NET data, information and knowledge to help achieve FEWS NET's mission to sustainably prevent food insecurity and famine. During Y1, the Hub conducted a series of technical reviews and stakeholder analyses to the development of its Y2 Annual Revision Plan for the Hub's DMP; this Gender Review and Gap Analysis provides the background to the gender-related findings.

The Hub's TO lays out the importance of gender considerations to all of the Hub's work. Before delving into the particulars of the scope, Section C.2 reads: "The capacity for FEWS NET to incorporate gender-sensitive analyses, metrics, and methodologies to the greatest extent feasible is one of the overarching principles guiding implementation of the Hub Task Order."

Section C.3 on Gender Considerations reflects the emphasis on the following:

> *Gender equality and female empowerment are core values of USAID policy and practice, and underscored by numerous Presidential Memoranda, Executive Orders, and national strategies directing the U.S. Government to ensure the meaningful involvement of women in foreign assistance efforts, and to address the community development and security challenges women, girls, and LGBTI people face. Understanding and adapting programming to address the different roles of males and females in access to, and use of, food security data, information systems, and food security-related assistance can lead to greater female empowerment and have a positive influence on food security and nutrition outcomes, especially at the local and national levels.*
>
> In some of the monitoring and assessment activities carried out in FEWS NET 7, gender will have direct relevance to food security analysis and decision-support activities.
>
> *The Contractor will provide sex-disaggregated data and analysis when gender differences are known and/or identified by USAID. The Contractor will use gender-sensitive methodologies in all relevant activities and analyses outlined in this SOW. The Contractor will describe in any relevant required reporting gender-based successes and challenges related to the food-security efforts, as outlined in this SOW. The Contractor may also be asked to improve or establish gender disaggregation in data and analysis when gender differences are unknown and unclear. The Contractor is expected to identify data gaps and gender analysis gaps, as well as to collect and disseminate gender information obtained or produced under this contract.* (pp 19-20)

The Hub TO 1 does not address how the Hub should relate to the other implementing partners. Presumably, as the Hub does not itself usually collect data, it will collaborate with other implementers to ensure that gender-relevant data are at least properly captured, archived, and disseminated. Depending on the extent to which other FEWS NET implementers have received similar instructions, the TO suggests that the Hub could proactively set up protocols to mainstream gender considerations into data collection and other research. For example, in describing the functions of the other pillars, the Hub's TO 1 suggests that Pillar 3 could play a key role in furnishing Gender information:

> *Pillar 3: Analysis of the dynamics of food, nutrition and livelihood security: Activities under Pillar 3 will deepen an understanding of the causes of persistent or recurrent food insecurity, hunger, malnutrition, vulnerability to food insecurity, and lack of resilience. This pillar will utilize and expand on FEWS NET's rich understanding of livelihoods, markets, agro- climatology, nutrition, and other physical and socio-economic phenomena, including gender and intra-household dynamics. It will apply this understanding to identify solutions which sustainably enhance food, nutrition, and livelihood security, and build resilience.* (Section C.1, p 9 of the Hub TO 1).

METHODOLOGY

As an initial step to incorporate USAID's vision for integrating gender considerations into the Hub's data-related activities, Hub staff:

1. Interviewed the Hub's Contracting Officer's Representative (COR) to determine the appropriate scope of this review;
2. Reviewed whether the FDW has, or might potentially incorporate, sex-dissaggregated data. Information was primarily provided by Roger Hunwicks, the DMP Technical Lead and FDW architect since 2012;
3. Reviewed the extent to which the FEWS NET website provides gender-related analyses, information or data; and
4. Interviewed the EW implementer on its current and anticipated use of such data. Some information was gathered from Data Management Advisor Nour Nourey, and Senior Markets and Trade Advisor, Sonja Perakis in interviews about the larger DMP AVP; responses to more specific questions about the EW team's use of, and plans for, sex-dissaggregated data were provided by Quality Assurance and Methods Advisor, Peter Thomas.

From a methodological perspective, this report will not delve into the intricacies of the vast literature defining the gender dimensions of data collection, handling, analysis, and reporting. For the purposes of this discussion, the term "gender" refers to roles and "sex" refers to biology. The vast majority of field-data from FEWS NET countries is likely to be relevant to the simpler biological distinction rather than the complex cultural continuum of roles and sexuality. Gender analyses are typically based on data that is sex-disaggregated but can build out that binary distinction to incorporate cultural dimensions. This review focused on identifying sources of sex-disaggregated data with the expectation that this distinction could ultimately provide a baseline for a greater gender analyses.

RESULTS OF THE REVIEW

SCOPE

After years of conceptualizing, designing, and procuring this phase of the FEWS NET project, USAID awarded contracts to two of what will ultimately be several implementers of the new, expanded, FEWS NET team. As laid out in two key presentations during the early months of the project,[26] USAID wholly re-conceived the role played by data and information. The Hub was created to serve as a knowledge management platform that serves not only the FEWS NET team (itself a coalition of several organizations) but also the greater food security and early warning community worldwide.

USAID expects the Hub to mature into this role over time, as earlier working relationships are recast and new ones forged. Due to particularities of contracting, the work plans for the two main contractors–the Hub and the EW implementers–were developed separately and are several months out of phase. One of USAID's early priorities has been to put in place a joint work planning process to align activities and timelines during the second year of the current phase of FEWS NET. In the meantime, USAID has taken initiatives at several levels within the limits of each implementer's current work plan to coordinate related activities through Sprint Planning Meetings, Team-Wide Collaboration meetings, as well as analyses of the "gray areas" that arise as work plans get implemented.

In this context, the COR proposed that this gender review focus primarily on the Hub's own portfolio. The issues thus identified can then be discussed at upcoming Team-Wide Collaboration meetings and incorporated into the joint work planning process currently scheduled for July-August 2020.

GENDER IN THE FDW

Table 1 illustrates that no sex-disagrregated data is present in the FDW. Nor is the Hub aware, from its Sprint processes, of any plans to load sex-disaggregated data into FDW.

Demographic data. There is a plan to upload sex-disaggrated population data. Currently, FDW provides total population estimates for approved geographic units, such as administrative units, food security classification units, and crop reporting units. As described in the Hub Population Data Strategy, the Hub will upload Landscan data, which will provide age and gender breakdowns for those geographic units, in Y2. This will allow analysts to determine the demography of geographic units tied to other data, like price and crop production, but will not provide sex-disaggregated indicator values for those series. Furthermore, due to changes in the Landscan methodology over time, it will not be possible to create time series analyses for the sex-disaggregated population counts.

Finally, the Hub does not currently plan to let users extract population data by drawing their own polygons or upload temporary shapefiles; to get demographic data, including sex-disaggregated data for arbitrary shapes, will require further research and considerable additional work.

[26] The Team-Wide Consultations on Information Flow and Coordination Report of March 12, 2020 and the FEWS NET Team Annual Meeting of April 24, 2020

Nutrition data. The nutrition domain in FDW captures indicator values for arbitrary indicators, for specific localities. At present, the FDW contains nutrition data only for Burkina Faso, Ethiopia and Mauritania; and the only values for each indicator value are household sample size, adult sample size and child sample size. If the original data, to which the Hub currently does not have access, did include gender information, the Hub could revise FDW's nutrition domain to capture the disaggregated values. Alternatively, the Hub is currently exploring the use of semi-structured formats, rather than rigid domain structures, as a means for archiving data sets that do not easily conform to a standard form.

Table 1. FEWS NET Indicators

DATA DOMAIN	NAME OF INDICATOR	SOURCE OF INDICATOR	PRESENCE OF SEX-DISAGGREGATED DATA
Crop Production	Area Harvested	FDE	no
	Area Planted	FDE	no
	Quantity Produced	FDE	no
	Yield	FDE	no
IPC Population	IPC Population Size	FDE	no
Nutrition	% GAM MUAC <12.5 cm	FDE	no
	% GAM WHZ <-2 2006 WHO Standards	FDE	no
	% SAM <11.5 cm	FDE	no
	% SAM WHZ <-3 2006 WHO Standards	FDE	no
	% Stunting HAZ <-2 2006 WHO Standards	FDE	no
	Anemia among children under 5 years %	FDE	no
	Anemia among pregnant women %	FDE	no
	Anemia among women of child bearing age %	FDE	no
	Crude Death Rate deaths/10,000 people/day	FDE	no
	Crude Mortality Rate	FDE	no
	Deworming among children under 5 years %	FDE	no
	Diarrhea %	FDE	no
	Iron-folic acid supplementation of pregnant women %	FDE	no
	Latrines %	FDE	no
	Malaria/ Fever %	FDE	no
	Measles %	FDE	no
	Measles vaccination among children under 5 years %	FDE	no
	Under 5 Death Rate deaths/10,000 children <5 years/day	FDE	no
	Under 5 Mortality Rate	FDE	no
	Upper Respiratory Infection %	FDE	no
	Vitamin A supplementation of children 6-59 months %	FDE	no
	Body Mass Index	FT: Scenario Devel.	no

DATA DOMAIN	NAME OF INDICATOR	SOURCE OF INDICATOR	PRESENCE OF SEX-DISAGGREGATED DATA
Population	Number of Households	FDE	no
	Population Calculated	FDE	no
	Population Census	FDE	no
	Population Estimate	FDE	no
Response	Amount of Food Commodity (MT)	FDE	no
	Amount of Food Purchased Locally (MT)	FDE	no
	Amount of Food Purchased Regionally (MT)	FDE	no
	Cash Transfer (USD)	FDE	no
	Number of Beneficiaries of Cash	FDE	no
	Number of Beneficiaries of Food Commodity	FDE	no
	Number of Beneficiaries of Food Vouchers	FDE	no
	Total Number of Beneficiaries	FDE	no
	Total Operation Cost (USD)	FDE	no
	Value of Cash (USD)	FDE	no
	Value of Food Commodity (USD)	FDE	no
	Value of Food Purchased Locally (USD)	FDE	no
	Value of Food Purchased Regionally (USD)	FDE	no
	Value of Food Vouchers (USD)	FDE	no
Trade		FDE	no
Price		FDE	no
IPC/FIC		FDE	no
IPC FIPE		FDE	no
Exchange Rates, CPI, CPR		FDE	no
Food consumption, energy intake	Dietary energy intake	EW/IPC	?
	HDDS	EW/IPC	?
	FCS	EW/IPC	?
	HHS	EW/IPC	?
	rCSI	EW/IPC	?
	HEA	EW/IPC	?
HEA Dashboard	Types of Hazards	EW	?

GENDER IN EARLY WARNING TEAM ANALYSES

Similar to the Hub, the EW team reports that their TO also underscores the importance of incorporating gender considerations into their data collection, handling and analyses. As for the Hub, this element is new to the project and currently under development; the EW team is exploring the links between vulnerability and gender.

According to Peter Thomas, the EW team currently accesses a limited amount of sex- and/or sex-disaggregated data in a systematic way across countries. The reason is largely two-fold: i) because broad responses to food insecurity tend to be funneled through humanitarian food assistance which targets households, not individuals, most data used in acute food security analysis does not describe individuals, but rather households or areas; and ii) constraints/opportunities related to sex and gender vary across contexts, and as such, relevant aspects to monitor related to the intersection between gender and acute food insecurity also vary significantly across contexts.

Nutrition data. Often, nutrition surveys contain some measure of sex-disaggregated data, typically to produce statistics for pregnant and lactating women. The surveys also include the sex of children being measured, although this breakdown is typically used to validate the strength of survey implementation rather than for publishing results by sex of child; such information is often used to infer the inter-household allocation of resources.

EW team Data Management Advisor Nour Nourey reports that in the course of uploading nutrition data to the DDL, she often spots data disaggregated by sex and age.

Livelihoods data. The information collected in support of FEWS NETS Livelihoods Analysis often includes sex-disaggregated information, particularly for female-headed households, when potential differences are identified during livelihood exercises. Other rapid assessments also identify potential differences for female-headed households (for example, among displaced populations). However, in these contexts, monitoring data streams is sporadic, with little to no consistent data streams including sex- and/or gender-disaggregated data.

Markets and trade data. Basic data on the supply and demand for agricultural commodities are not, per se, associated with gender. However, Perakis of the EW team noted that organizations such as WFP, incorporate gender into their analyses of vulnerability to market-related shocks.

Monitoring and Evaluation data. The EW team typically breaks out its training data by the sex of the participants.

GENDER ON THE FEWS NET WEBSITE

At present, a search of the FEWS NET website fails to display any documents related to the terms gender, women, woman, or female. Those terms may, however, exist inside PDF reports, since the current search engine only scans report titles, not content. This result is not surprising considering that content for the website is provided by the EW team, which does not currently systemically incorporate gender into its analyses or data.

GENDER IN THE HUB-MANAGED DOCUMENT ARCHIVES

As part of the transition from the prior iteration of the project, USAID tasked the EW team with sharing relevant project documents with new team members, including the Hub. The documents, which represent an Archive of past documentation and thus will not be updated, are now housed by the Hub on the FEWS NET Exchange, the intranet site for the Project.

A cursory search conducted in the Google Drive folder of the Exchange found that of the 200+ documents in the Archive, only *five* documents mentioned gender, women, woman, or female (Table 2). Of those five documents, only *three* represented substantive content related to gender; the remaining two documents only referenced gender using descriptive terminology and were thus deemed irrelevant to this review.

While the EW team may hold additional reference materials, the Hub-run archive does not currently hold significant gender-related resources.

Table 2. References to gender in the document archives

REFERENCE DOCUMENT AND SEARCH TERM(S)	YEAR	RELEVANT TEXT
Enhanced Market Analysis Report_Product Documentation (Gender)	2019	Gender dynamics and behavior in market contexts to inform market analysis reporting
Uganda Network Development Strategy_Multi-Hazard Bulletin (Gender)	2017	N/A, descriptive
FSIN Operational Strategy (Gender, Woman, Women)	2016	N/A, descriptive
Network Development Strategy (Woman, Women)	2013	Demographic details of FEWS NET training participants
FDW Nutrition Data Domain_Process Documentation (Woman, Women)	N/A	Documentation for 3 nutrition data domains (iron supplementation and anemia among pregnant women)
No references to Female or Sex		

DISCUSSION

The reason for FEWS NET's irregular inclusion of gender considerations is largely two-fold: 1) the great majority of food security interventions targets areas or households, not individuals, and therefore the information on needs, distributions and responses reflects those same units; and 2) the constraints/opportunities related to sex and gender vary across contexts, and thus the relevance and form of the data needed to monitor the links between sex and acute food insecurity also vary significantly across contexts.

Other organizations are incorporating gender into their analyses and knowledge products. Not only is Gender Equality one of the United Nations' 17 Sustainable Development Goals, but women are specifically highlighted in the targets - and thus indicators - for several other goals, including ending hunger.[27]

The International Food Policy Research Institute has conducted decades of research on the interactions between gender, food security and nutrition (https://www.ifpri.org/topic/gender). Some of that work served as the basis for the development of the Women's Empowerment in Agriculture Index, which has for years been a formal indicator in the Feed-the-Future Monitoring System that underlies the USAID's Feed the Future program. In the humanitarian field, organizations like WFP, CARE and Catholic Relief Services, which use their food assistance programs to empower women, will often analyze women's roles and vulnerability to shocks in related marketing systems.

FAO in its March 2019 evaluation of the IPC noted in the section on Equity and Gender that "A number of IPC users interviewed for this evaluation raised concerns about the lack of disaggregation in the IPC [Acute Food Insecurity Scale] analysis. The greatest demand is for disaggregation to smaller geographical units, followed by disaggregation by different population groups. Addressing the latter is planned in the next phase of the Global Strategic Program."[28]

Within FEWS NET, efforts are underway to improve the project's capacity to support gender analyses. While the existing FDW domains do not yet have sex-dissaggregated indicators defined, these could easily be incorporated when such data becomes available. At the end of Y1, the Hub was working to incorporate demographic data from LandScan into the FDW, including the number of males and females at the first administrative level in each country.[29] As part of that plan, the Hub is testing how these data might be applied to FEWS NET's other predefined units (e.g., livelihood zones). Further, the Hub could add additional fields to the nutrition domain, or, at a minimum, store sex-dissaggregated data regardless of formats in the semi-structured domains currently under development.

[27] United Nations Women SDG 2: End hunger, achieve food security and improved nutrition and promote sustainable agriculture. Available online.
[28] Food and Agriculture Organization of the United Nations (2019) Final Evaluation of the Integrated Food Security Phase Classification (IPC) Global Strategic Program 2014-2018. Page 5. Available online.
[29] For example States in the US, or Provinces in Afghanistan, Canada or Zambia.

CONCLUSIONS AND RECOMMENDATIONS

The Y1 Gender Analysis conducted by the Hub serves as an initial stock-taking effort for FEWS NET. To meet the gender objectives laid out for this phase of the project, the Hub must work with other FEWS NET team members to explore how to add useful, sex-dissggregated data and tools to food security analysts inside and beyond the project. Possible next steps may include:

- Working with the EW team to do a more thorough analysis of what gender-related data and analyses are being used both by their own analysts as well as the panoply of organizations that partner with FEWS NET at the national, regional and international level.
- Reviewing tools and methods for handling sex-disaggregated data; the FDW and Data Portal need to support a consistent method for disaggregation and analysis by gender and age-group for all relevant household and individual datasets. The developer Organisation for Economic Cooperation and Development Gender Data Portal could carry out an external consultancy to this end.
- Designing data and archiving protocols to support gender analysis, which would ensure that, where relevant, all datasets are coded to allow consistent gender and age disaggregation, and that query and other tools are similarly designed for disaggregated results.
- Incorporate research on the role of gender analyses in improving early warning and food assistance targeting under Pillar 3 (described above).
- Identifying the expertise, assets and roles within the FEWS NET team related to gender analysis. One approach is to use existing collaboration platforms, such as the monthly Team-wide Collaboration meetings, the 2020 Joint Work Planning process, and future FEWS NET Forums as an opportunity for the FEWS NET team to address such issues.

ANNEX 6. FEWS NET COMPLIANCE REVIEW

The following report was commissioned by the Hub's external partner Insight Systems. The report below represents an edited version of their final deliverable submitted to the Hub, presented below in its entirety. Any references to "we" refer to Insight Systems.

PURPOSE

The purpose of this project is to assess the FEWS NET Learning and Data Hub security boundary definition to determine the security actions needed for the FEWS NET system to bring FEWS NET into compliance and obtain its ATO from USAID. The process will include formal boundary definition, security categorization of the system, and high-level costs for both the SA&A preparation review and assessment.

IN SCOPE

The scope of this project is to create a formal boundary definition and categorization of the FEWS NET "System," which comprises the FDW, FDW, and the FEWS NET website (fews.net).

The above components will be reviewed during this assessment. Based on the system boundary determination, a security categorization review will be conducted to determine the work required to prepare for an SA&A and high-level costs for the preparation support and assessment. This assessment is based on the National Institute of Standards (NIST) 800-53 Revision 4 which is the current version being used by USAID. The deliverables are based on information provided by Kimetrica.

OUT OF SCOPE

The following items are out of scope for this assessment:

- This effort will not prepare for the security assessment and will not provide any security assessment documentation required for the system authority to operate at USAID; and
- This effort is limited to the FEWS NET "System" as defined above. No other components will be evaluated.

DELIVERABLES

The deliverables of the Hub Assessment are provided in sections below.

BOUNDARY CATEGORIZATION

The first deliverable of Hub's Assessment is the System Categorization document which is the USAID standard template for the NIST Federal Information Processing Standards Publication 199 (FIPS 199). The document was created based on information provided by Kimetrica. The security categorization was determined to be "Moderate." A third-party assessment must be procured for systems categorized as "Moderate" per the USAID Automated Directive Service (ADS) 545 (USAID Information

Technology (IT) Policy). The FEWS NET Hub Categorization document has been provided as a separate MS Word document called "FEWS NET Hub System Categorization."

In addition, M/CIO/IA added the requirements to submit an Information Systems Inventory (Form 502-1) for system data retention. Insight provided the bulk of this additional document to Kimetrica (minus reference documentation information). The document is called "FEWS NET HUB AID 502-1 (Information Systems Inventory)." Kimetrica can review, add the documentation information, and submit the document to recordsinquiry@usaid.gov, at its discretion.

Note: The Categorization was created based on the website being connected to the Data Warehouse.

APPLICABLE SECURITY CONTROLS

Based on the Categorization, Insight put together the list of applicable controls and provided common control sets in which the FEWS NET Hub reviewed for inheritance. This included the Agency, CIO, information assurance (IA), and Office of Security, Privacy and AWS Few Ramp package common controls. The USAID Common Control sets for the Agency, CIO, IA (which are currently being consolidated) and Office of Security are provided in an excel spreadsheet called "USAID Master CIS_Rev 1.4." The Master Center for Internet Security common controls can be applicable to the system as they define policies and activities that USAID defines for its systems. Requirements for inheritance are provided within the document by security control.

The Insight team also reviewed the AWS East/West Fed Ramp package and provided a list of Infrastructure as a Service and Platform as a Service[30] controls available for inheritance. The list of potential AWS common controls can be found in the "FEWS NET AWS and Privacy Common Controls 05062020" document. In addition, the USAID Privacy common controls were elaborated on in the "Privacy Common Controls SSP 2020048" document.

SUPPORTING DOCUMENTATION

Based on the controls that are required for the FEWS NET Hub security assessment, the Insight team pulled the templates of the required documentation and placed them in a zip file (called SA&A Documents) for the FEWS NET team to access. **Note:** The USAID Information Assurance branch (M/CIO/IA) does make periodic changes to its templates. These changes are usually not significant but these documents represent a point in time. A recommendation for the order of documentation completion is provided in Table 2 of this report.

In addition, the team has furnished a sample list of evidence that is usually requested as part of the assessment process in the document titled "Sample Security Controls Artifacts." It is useful to attach this evidence to the controls as the package is being put together in the USAID security package tool, Cyber Security Assessment and Management (CSAM). **Note:** this list contains artifacts by control. If a control is not applicable, the evidence associated with it is not required.

[30] Commonly referred to as IaaS and PaaS, respectively.

COSTS OF SECURITY ASSESSMENT AND AUTHORIZATION

The costs for third party Security Assessment and Authorization are published on the USAID Intranet on the USAID's M/CIO/IA page. These costs are presented in Table 1.

Table 1: Costs of SA&A, by Size (XS to XL)

		XS	S	M	L	XL
SCOPE	Physical/Virtual Servers/ Nodes: To include Dev, Test, Prod	1-6	1-6	7-25	26-50	51+
	Applications/Subsystems	1	1-2	3-4	5-6	7+
	FIPS Categorization	Low	Moderate	Moderate	Moderate or High	Moderate or High
	Customer Responsibility Controls only leveraging FEDRAMP[31] SA&A	25	25-50	50-150	50-150	50-150
1/3 ASSESSMENT		$20,000	$27,500	$30,000	$40,000	Contact IA/CRM
FULL SA&A		$29,130	$38,841	$58,261	$72,826	Contact IA/CRM
Customer Relationship Manager (CRM) ONLY		$14,423	$19,231	$28,846	$43,269	Contact IA/CRM
Services included are SA&A Status Reports (distributed weekly to System POC's during the project); and Assessment Out-brief(s) with the Authorization Office and Security Office. **Note:** SA&A cost may vary from the fixed prices stated here to adjust for level of effort fluctuation as a result of inherited control determination.						

[31] The Federal Risk and Authorization Management Program (FedRAMP) is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services.

DELIVERABLES INCLUDE:

Documents Created by assessor (Full SA&A)

- Out-brief presentation with findings summary
- Certification Letter
- ATO Letter, Security Assessment Plan (SAP)
- Security Assessment Report (SAR), including embedded Risk Assessment Report (RAR)
- Plan of Action & Milestones (POA&M) with Mitigation Recommendations

Documents created by Assessor (1/3 assessment)

- Out-brief presentation with findings summary
- Certification Letter
- SAP
- SAR, to include embedded RAR
- POA&M with Mitigation Recommendations

Documents created by Assessor (FEDRAMP CRM)

- RAR based on the review of FEDRAMP package
- Out-brief presentation with findings summary
- Certification Letter
- ATO Letter
- SAP
- SAR, to include embedded RAR
- POA&M with Mitigation Recommendations

As demonstrated in Table 1, M/CIO/IA is implementing SA&A as a service to provide basic input/output systems with an independent assessment service for a fixed cost that is based on a simple formula to categorize systems based on size and type of assessment required. Cost of Phase 1 documentation development (e.g., System Security Plan and other supporting documentation etc.) and Information System Security Officer (ISSO) support services are not included in the fixed cost annotated in the table. The Hub should plan to budget these costs separately and work with the M/CIO ISSO team to develop individual cost estimates.

Size LARGE determination. Based on the Categorization of the system, the size determination was "Large." The cost for the initial full third-party assessment will be $72,826. This estimate was verified by the Assessment Lead, Michele Brantley-Cotton, in an email. The verification email, titled "FEWS NET Assessment Cost," has been included in the deliverable package. The current estimate for subsequent yearly 1/3 assessments (required for continuous monitoring of the system) is currently $40,000 per year. **Note**: USAID is currently on NIST 800-53 Revision 4. However, NIST 800-53 Revision 5 (Rev 5) should be available within the next year. This will increase all assessment costs. An additional $10,000 per year should be estimated for annual 1/3 control review continuous monitoring costs in the future. Also, once a system is fully assessed, only 1/3 control review is required moving forward if the system consistently follows continuous monitoring.

USAID M/CIO/IA does offer ISSO services for systems at a cost. These services do not include the initial assessment but do include annual assessments of the system. This may be a viable option as some of the annual costs will be covered under the ISSO services. If interested in procuring these services, contact Dmitriy Radchenko in M/CIO (dradchenko@usaid.gov).

SA&A PREPARATION REVIEW AND SUPPORT COSTS

The Insight team can assist the Hub in preparing for assessment, and provide assessment support, including the following:

- Review and feedback on supporting documentation;
- Review and feedback on System Security Plan Control Implementation statements;

- Updating the CSAM package with System Security Plan implementation statements and associated documents provided by Kimetrica;
- Generating the System Security Plan based on information given, and providing it for signature;
- Supporting the team through security assessment to ensure that the process is followed and that no undue findings are associated with the system; and
- Work with the team to define a plan of action and milestones for any findings resulting from the assessment.

The SA&A support effort will be approximately 80 hours at a cost of approximately $16,000.00. The team will not act as an ISSO for the system but can assist the team in reviewing and liaising with M/CIO/IA. It is recommended that if the team does not have a dedicated ISSO, that the ISSO services be procured from M/CIO/IA. With this option, no additional support is required.

Table 5. Documentation Order Recommendations

DOCUMENT NAME	COMMENTS
Categorization	Complete – Please review, sign and route to CRM for signatures.
502-1	Completed first pass of this document. Please review, update and submit to recordsinquiry@usaid.gov.
Technical Architecture Document	The Technical Architecture Document must be delivered to M/CIO Engineering Review Board/Architecture Review Board for review and approval. This process will also take the system through becoming an official USAID-approved application.
System Owner Designation Letter	Voluntary Product Accessibility Templates information is not required as the system was built before January 18, 2018.
ISSO Designation Letter	Categorization, 502-1, Technical Architecture, System Owner Designation and ISSO Designation are all required to begin the SA&A process and create a package in the security authorization system, CSAM.
Business Impact Assessment	Will need to define the Recovery Point Objective, Recovery Time Objective, and Maximum Tolerable Downtime.
Privacy Threshold Assessment (PTA)	No Personally Identifiable Information (PII), only PTA should be required.
System Boundary Spreadsheet	
Technical Description Document	This is uploaded into CSAM and is generated as part of the System Security Plan.
System Description Document	This is uploaded into CSAM and is generated as part of the System Security Plan.
Configuration Management Plan	
Change Management Plan	
Contingency Plan	
Incident Response Plan	
Contingency/Incident Response Test Results	It is beneficial to merge these tests into one test results document as they are closely associated and usually require the same resources. Must record all participants in the test. Should be representative of all roles in the plans.
e-Auth Risk Assessment	

Access and Identity Management Plan (Include Separation of Duties)	These are currently three templates, but they are closely aligned. It is advisable to merge these components into one document.
Audit and Accountability Procedures	
Systems Security Plan	Please see the list of controls. Implementation statements for all applicable NIST 800-53 Rev 4 controls are required.
Maintenance Procedures (MA-1)	This may be addressed in the MA-1 control implementation statement. Document is not required.
Media Handling Procedures (MP-1)	This is most likely not required. It can be documented in the MP-1 control.
System and Communications Protection Procedures (SC-1)	This may be addressed in the SC-1 control implementation statement. Document is not required.
System and Information Integrity Procedures (SI-1)	This may be addressed in the SI-1 control implementation statement. Document is not required.
Physical and Environmental Protection Procedures	Physical and Environmental Controls should be fully inherited from the Fed Ramp package.
Security Assessment & Authorization Policy and Procedures (CA-1)	
Operations and Maintenance Plan	
Continuous Monitoring Schedule	This is required once the system is authorized. Must be submitted monthly (Friday of the first full week of each month).

ANNEX 7. HUB BACKLOG REVIEW

TYPE	KEY	SUMMARY
Task	DATA-50	Investigate a Python only approach to boundary transfer/derive a boundary for a Geographic Unit from its Relationships
Task	DATA-71	Uploading Xborder Trade data
Task	DATA-81	Establish and implement procedures for data quality assurance (SOP format)
Task	DATA-82	Develop a Data and Data System Documentation and Training Strategy
Task	DATA-89	Conduct Data Domains Review
Task	DATA-95	Run FDE Tests as part of FDW continuous integration
Task	DATA-105	Visualization Tools Strategy
Task	DATA-116	Produce a data set development strategy
Task	DATA-117	Develop strategy to establish linkages with the data sets and tools from FN Team and other partners
Task	DATA-185	Test the Cloudfront approach
Task	DATA-193	Science Team gets FEWS NET access level: how do they know what data are sensitive?
Task	DATA-206	Missing LandScan Country Population records
Story	DATA-60	Prototype: API Linkage to Democratic Republic of the Congo Price Data
Story	DATA-61	Prototype: API Linkage to FN Price Data Collection Tool
Story	DATA-62	Prototype: Automated Online Price Bulletin Report
Story	DATA-92	Y1 Extensive Technical Review of Data Sets
Story	DATA-113	Develop a Data Ownership, Data Permissions and Data Sharing Strategy for National Data Source Organizations
Story	DATA-114	Develop data sharing/collaboration strategy & protocols for other non-Team users
Story	DATA-115	Develop an approach to international coordination of data quality improvement, data-sharing, and data harmonization (such as FSIN, FAO, WFP, IPC, World Bank)
Story	DATA-126	Allow user control of notifications
Initiative	DATA-209	Data as a product: Thinking about Crop Calendars and more, in the FDW
Bug	DATA-45	String too long to represent as jsonb string
Bug	DATA-73	The aliases are dropped off at the creation of the product codes
Bug	DATA-74	Editing common products codes generates 500 Internal server error
Bug	DATA-125	Dates on slider do not correspond to price dates
Bug	DATA-203	It is the Not Found error message which is shown upon trying to open Documents in a new tab of the Mayan system

ANNEX 8. STAKEHOLDER CONSULTATIONS

BACKGROUND

The FEWS NET Learning and Data Hub conducted a series of stakeholder consultations to support the technical analyses for its *Annual Plan for Revisions* to the FEWS NET DMP. The Hub conducted these stakeholder consultations during March and April 2020 with both generators and users of FEWS NET data. Within the FEWS NET team, the Hub interviewed the Pillar 1 implementer, the EW team, and the Pillar 3 Program Manager. Following identification of stakeholders by the COR team as regular users of FEWS NET products, the Hub consulted representatives the USAID BHA—Office of United States Foreign Disaster Assistance and Disaster Data, Assessments, Technology & Analysis (DDATA)—and the BRFS (including the Food Security Monitoring and Evaluation Group, Analytics and Learning Division, and Resilience teams).

To improve policy and decision-making, discussion focused on how respondents use the FDW and the data within it. These discussions covered a range of topics, and provided the Hub with a number of insights into how the Hub might improve its DMP. All interviews were conducted virtually.

RESULTS OF EARLY WARNING TEAM CONSULTATIONS

The EW and Hub teams frequently coordinate on data management, storage, archive, and analytical archiving activities related to the FDW and FDE. Regular collaboration includes the monthly Data Stakeholders Meeting convened by the EW team and the monthly Sprint Demo and Sprint Planning meetings, organized by the Hub DMP. These meetings generate a list of data-related tasks and services for primary management by the Hub, as prioritized by USAID.

To engage the EW team in strategic reflection, the Hub held several meetings in April 2020 with EW headquarters (HQ) staff to identify how the Hub's DMP can better serve their data requirements. The Hub identified the interviewees specifically for the different ways they relate to the Hub's DMP. Discussion focused on each interviewee's differing uses and needs for FEWS NET data. While they were interviewed separately to provide time for each to offer in-depth comments, their remarks have been consolidated below by data set or analytical process to facilitate integration into the Hub's work planning exercises. The following section includes the perspectives of:

- ***Nour Nourey,*** interviewed on April 10, 2020. As the EW team's Data Management Advisor, she manages activities related to data, the website, GIS, information technology, and network support. She also serves as the point-of-contact for the Hub on these activities. Nourey frequently reviews, cleans, manipulates, and organizes data provided by FEWS NET field analysts, and less frequently, for HQ analysts who rely on her to format the data they require from the FDW.
- ***Peter Thomas,*** interviewed on April 14, 2020. As the EW team's Quality Assurance and Methods Advisor, he identifies ways to improve the EW team's analyses, including exploring

new tools and leading retrospective evaluations of FEWS NET's analyses on both a regular and ad hoc basis. His primary data-related activities include i) supporting early warning classification and analysis; ii) accessing complementary external surveys and baseline data from Department of Homeland Security, Food for Peace (FFP), and Feed-the-Future; and iii) working with sector and outcome data from the EW team's integrated food security analysis framework.

- ***Sonja Perakis,*** interviewed on April 15, 2020. As the EW team's Senior Markets and Trade Advisor, she provides overall technical oversight and guidance related to market analysis, including both market monitoring and market projections. Her role includes data collection, processing, and analysis.

FEWS NET DATA SETS

Questions regarding the data used to carry out the EW team's responsibilities revealed a division between data sets currently included in the FDW/FDE and those that are not.

DATA INCLUDED IN FDW/FDE

The following sections cover data sets that are currently included in the FDW/FDE in some form.

Consumer Price Index data. These data are sourced from national Ministries of Finance and uploaded into FDW by field staff.

Crop data. Nourey's team currently prepares these data on an annual basis; however, she proposes that field staff could handle this task, shifting LOE. Current challenges include:

- ***Timeliness of sub-national estimates.*** In cases where data are delayed or unreliable, the EW team use either their expert judgment or remote sensing methods to derive subnational production estimates (e.g., food insecurity estimates may reflect subnational crop production estimates derived from the Water Requirements Satisfaction Index).
- ***Accommodating disparate formats.*** Because the data arrive in Microsoft Excel spreadsheets and Word documents, the EW team must spend significant LOE to translate the data into spreadsheets using Tabula, a specialized software.
- ***Converting units.*** The data often need to be converted from local measures into metric.
- ***Updating the data.*** The EW team will maintain updated data for FEWS NET presence countries, but is unable to keep the other 140 countries in the database up-to-date.
- ***Standardizing the metadata.*** Maintenance of frequently changing FEWS NET Identifiers (FNID) for geographic entities and source document names is labor intensive.
- ***Data cleaning.*** Field staff offices differ in how they provide data and metadata, varying from country to country. As time continues, the data will require updating through manual or automated data cleaning.

Exchange rates data. The Markets and Trade team uses international market exchange rates, downloaded from OANDA Corporation and loaded into FDW by the EW Home Office data team, as well as official government rates (published by the relevant government) and secondary exchange rates (i.e., parallel or black market rates). Government and secondary exchange rates are obtained by FEWS NET's field staff and then compiled and uploaded into the FDW by the Field office staff.

Food availability data. The EW team creates commodity balance estimates using production, consumption, and trade data, which are provided by secondary sources' seasonal assessment reports. The EW team typically relies upon draft reports, which provide projections rather than final estimates and may or may not reflect the data later uploaded into the FDW.

Humanitarian responses data. At present, these data in the FDW are dichotomous (yes/no) and infrequently updated, limiting its utility for the EW team (e.g., USAID frequently requests analyses on how food assistance impacts markets) and requiring analysts to maintain their own versions of these data. The EW team recommends identifying solutions to more easily extract current data at national and subnational levels, ultimately allowing for smoother integration with other FDW data.

Integrated Phase Classification data. The EW team forecasts emergency food assistance needs in each FEWS NET presence country by comparing estimates of the acutely food insecure population to the previous year and the recent five-year average. The population is assigned an integrated phase classification Food Insecurity Classification. In the Food Security Outlook Briefs, this population with an IPC category is published as a range (i.e., "bin"); within FDW, Nourey uploads the actual estimates not the ranges. Upon extraction, FDW users external to the project receive the bin, while internal team members receive both the estimate and bin.

The EW team generally does not use the FDW version of these data, which are uploaded as national totals. For analyzing the poorer wealth groups, they require greater granularity .

As a core FEWS NET product, these data are in high demand and would likely receive broader use if more accessible upon extraction from the FDW. At present, these data require considerable manipulations to be useful for analysts. Specific improvements may include:

- Converting the projections from the current format to ML1 and ML2, HA1, HA2[32];
- Providing tabular data by IPC phase, broken down by 0/1 and 1, 2, 3, 4, 5; and
- Supporting spatial time series analyses.

Price data. The majority (80%) of FDW price data reflect national market information systems, including WFP, FAO and other partners, but in some countries (20%), FEWS NET field staff source these data. For the latter, field staff clean and aggregate the data before uploading into the FDW during the last two weeks of each month. To conduct their analyses, EW HQ analysts typically use the refreshable spreadsheets to retrieve price data from the FDW.

Population data. The EW team reported using numerous sources for population data, including the Landscan data uploaded by the Hub to FDW. However, analysts that need disaggregated subnational crop production data to estimate food requirements often use external sources, including country Ministries, the World Bank, and the South African Development Community. To make population data in the FDW more useful, changes to the following extraction processes were discussed:

[32] The EW team uses the following codes for their different IPC forecasts: CS for the Current Situation; ML1 for the Short Term, Most Likely scenario which covers the next 3 months; ML2 for the Medium Term, Most Likely scenario, which covers the 4-6 month projection; and HA1 and HA2 are the same as ML1 and ML2 but take into account likely food assistance.

- ***Livelihood zones, rather than administrative units.*** At present, livelihood zone shapefiles are not yet available. As such, the EW team extracts data at the lowest possible administrative level, then expends significant LOE to aggregate and match the data to livelihood zones.
- ***Country, rather than regional levels.*** At present, some population data in the FDW are available only by region (e.g., the Dominican Republic is only available as Central America). Providing country level data—as well as data on FEWS NET "watch countries," or those adjacent to presence countries and the broader set of metadata—would be beneficial.
- ***Demographic breakdowns.*** The EW team often requires demographic disaggregations (i.e., gender, age, or population densities for nutrition survey analysis or specific USAID ad hoc requests), which are not readily available in FDW at this time.
- ***Transpose capabilities.*** The EW team expressed interest in having the capability to transpose the data in the FDW, rather than manually following extraction.

Nutrition survey data. At present, the FDW does not contain many nutrition surveys, due in part to the LOE required for Nourey to extract these data from non-standard surveys. Additionally, because country Ministries often require authorization for data use, the EW team reported difficulties in distinguishing the appropriate FDW permissions level for these data. Furthermore, because the case for these data is not yet well-established, the Hub team may discuss data needs with EW team nutrition specialists and Decision Support Group analysts before revising this domain structure.

DATA NOT INCLUDED IN FDW/FDE

The following sections cover data sets that are currently not included in the FDW/FDE.

Afghanistan team data. Data from the Afghanistan office are primarily qualitative and generally include interviews conducted via Google Forms. Alternative approaches include tools such as Magpi and Kobo that offer table and browser interfaces to let users to define forms using spreadsheets or with drag/drop into lists.

Carry over stocks. To provide food balance estimates, the EW team maintains data on stocks carried over from the previous marketing year.

Conflict data. The EW team reported using data from Armed Conflict Location & Event Data Project (ACLED), as well as other sources for Yemen and Northeast Nigeria. The EW team downloads data from the ACLED website, manipulates it, and maps the data, often in ranges. Because the EW team only requires a few data points from ACLED, the LOE required is likely too high to justify uploading current ACLED data into the FDW. However, data that facilitates historic analyses of conflict may provide insight into local drivers of significant conflict events. Because ACLED's regular website updates overwrites earlier data, the FDW may be able to add value by capturing this historic data.

Gender disaggregated data. FEWS NET only recently began to incorporate gender into its methodology for food security analysis. During previous iterations of the project, the EW team engaged with the United Nations (UN) and Permanent Interstate Committee for Drought Control in the Sahel on including gender in early warning analyses. At present, the EW team is engaged with external partners to approach gender using vulnerability or capacities approaches. While no gender disaggregated data are currently in the FDW, these data are present in some nutrition surveys. Other

organizations in the food insecurity field, including WFP and Catholic Relief Services, maintain specific objectives to empower women in their programming (e.g., WFP analyzes the role of women in the marketing chain to identify differential shock impacts and associated food assistance needs).

Household survey data. The EW team gathers survey data informally—from other organizations with whom field staff have relationships—or formally through international sources (e.g., USAID's Demographic and Health Surveys and World Bank's Living Standards Measurement Surveys). With the former, field staff share data with regional teams who then clean, manipulate, and analyze the data, without modifying the original data. With the latter, HQ usually presents their analyses as Microsoft PowerPoint presentations, but does not share the survey data back with the field offices.

Because most surveys used by the EW team were originally conducted by organizations that do not seek to publish their results, data sharing agreements would be critical to upload household surveys into the FDW—none of which exist at the current time. Cleaning and processing instructions would also be necessary, since variable names and codes are often irregular in these informal surveys.

The EW team expressed interest in automating how household survey data are handled.

Internally Displaced Persons data. At present, the EW team removes Internally Displaced Persons (IDP) camp population estimates from the surrounding administrative area and classifies the two entities separately.[33] The EW team indicated that identifying IPC classifications for IDP camps would likely be beneficial, which are separate data points within administrative areas.

Livelihoods data. The FDW must provide clarity on the attributes of the livelihood zones to distinguish them from production systems. Furthermore, livelihood zones data in the FDW are outdated—some more than 10 years old—and must be updated to incorporate important structural changes. At present, the EW team is partnering with external organizations such as Save the Children or Oxfam for zoning changes and associated revisions to baseline estimates.

Per capita food requirements. The EW team multiplies per capita food requirements by population data to derive total food requirements, data are not currently in the FDW.

[33] There are no particular sensitivities with inputting IDP camp population estimates into the FDW.

Primary data. FEWS NET mostly uses secondary data (80%), with the exception of primary price data (20%). When authorized by USAID, the EW team will collect primary data to fill information gaps in the local humanitarian space (e.g., nutritional analyses in South Sudan and Ethiopia). Typically, the EW HQ designs the surveys, while field offices handle local logistics and hire local consultants to deploy the survey via tablets. The EW team stores these surveys on their server and submits them to USAID's Development Data Library. Field staff often keep a separate copy or receive a mirrored shared link.

Remittances data. The EW team monitors remittances data in Central America and Haiti, but it is unavailable for other FEWS NET presence countries. Remittances data for Guatemala is obtained from Banco de Guatemala and is loaded into FDW using the new Semi-Structured Domain by the Field Office. Remittances data for other Latin American and Caribbean countries may be loaded in the near future. The EW team currently uses proxy measures for remittances data, including international cash transfers.

Source documents. The EW team's database that stores source documents is extremely sensitive (e.g., when the EW team wants to change the standardized names of a source document, it can cause errors and/or will not allow information to be uploaded unless the name is an exact match).

Wage data. Though wage data are integral to non-farm incomes earned by both rural and urban populations, these data are unavailable in many locations. In locations where the EW team conducts market price surveys, they can identify local wages but lack ways to estimate the number of days the wage was earned. The EW team also estimates wages for casual workers by relying on key informants in vegetable, fruit, coffee, and construction (e.g., in Central America and Afghanistan).

Trade data. Currently, FDW contains the initial load of historic trade data for Southern Africa from 2015. Work is ongoing between the EW team and the Hub to reload the Trade domain with up-to-date data covering Southern and Eastern Africa. The EW team uses formal trade flow data, sourced from national governments, for countries that have well-developed trade monitoring systems (e.g., South Africa and Zambia). For most other countries that lack timely trade data for staple foods, the EW team stitches together mirrored export data from Comtrade.

The EW Markets and Trade team reported that they use data coming from sources outside the FDW to assess market functioning along key trade corridors.

Wheat bag symbols on IPC maps. IPC changed symbology from an exclamation point to a wheat bag symbol to indicate that a specific food security classification unit receives a significant amount of food aid.[34] However, the wheat bag symbol is not in the FDW, because these IPC classifications are not consistent with FEWS NET data, which continues to use an exclamation point to indicate that food security would be one IPC phase worse without current humanitarian food assistance delivery.

ANALYTICAL METHODS AND TOOLS

The EW team described key processes to handle, analyze, and present data.

[34] Note that use of the exclamation point protocol under IPC v2 officially changed to wheat bags under IPC v3.

Tools for storing, cleaning, manipulating, and analyzing data. The EW team reported that they plan to analyze the steps in report production. After convening an internal workshop to review the results, they will identify ways to facilitate the analytical process. At present, the EW team reported using the following tools:

- ***Data storage.*** While conducting analyses, each regional EW specialist stores a version of the data on a drive accessible on their computer, shared internally with the EW team. After publication, the data moves to a central EW drive. The team could consider using more extensive regional drives to facilitate common access (especially by new staff) and document their food balance computations for comparison with those from the FAO.
- ***Data analysis.*** The EW team mostly uses Microsoft Excel to calculate indicators and maintain records of the sources to protect against errors. They also use refreshable spreadsheets from the FDW to prepare tables and visualizations (e.g., price projections), which only requires manual data entry and automates the rest. The EW team uses SPSS to conduct retrospective analyses with legacy syntax that ensure consistent names and indicator ranges, as well as STATA, which is primarily used by the EW Markets team.
- ***Mapping.*** Following the Hub's assignment to add border points to an FNID in the FDW, the EW team reported that they separate market data from the corresponding geographic unit, download it in tabular format, and then conduct analyses and create maps in ArcGIS.
- ***Other visualization tools.*** Some members of the EW team use Tableau for visualizations, either online or through MS Exchange. They are interested in cloud-based applications that might be useful both for data exploration as well as creating the high quality visualizations required for the web.

Country Metadata Books. The idea of developing Country Metadata books, which would contain comprehensive metadata descriptions for all items used in each country, was introduced during the last phase of the FEWS NET project. This work, which was discussed and placed in the backlog prior to FEWS NET 7, may be reintroduced for discussion in the EW-led Data Stakeholders group.

Evaluating food security forecasts. The EW team uses two methods to do retrospective analyses. First, they compare the ML2 projection to the subsequent ML1 analysis to see if they are consistent. Second, they use survey data that comes out after-the-fact to generate food security outcome indicators which they then compare to their original analysis.

Training and documentation. At present, the EW HQ team conducts the majority of data cleaning, analysis, and presentation. If some of these steps were automated and staff had greater training in using the FDW, the EW HQ could shift some of this responsibility and LOE to field staff. This may be especially fruitful in analyses that require a sensitive yet strong understanding of local conditions (e.g., price forecasting). Guidance materials for the FDW could also provide information on the reliability and accuracy of the data.

DMP OPPORTUNITIES, AS IDENTIFIED BY THE EW TEAM

The EW team identified several ways the Hub could assist with their data-related needs, as well as the overall functioning of the FDW.

Data acquisition and upload. The EW team recommended several types of data for inclusion in the FDW. First, data not in the FDW but centrally organized and frequently used by the EW team to conduct food security analyses, including:[35]

- Real time data on changes in household income;
- Remittances;
- Wage data;
- Humanitarian response data (food assistance) by quantity at national and subnational levels;
- New livelihood baseline data; and
- Interannual social protection data at a subnational level.

Second, data the EW team manually curates, maintains, and combines using spreadsheets, which may or may not be centrally organized. The EW team suggested that the Hub work with both HQ and field staff to make an inventory of the full range of data sets and processes that might be best organized, stored, and available in the FDW. The inventory could identify which data acquisition and analysis processes could be automated; it would allow the project to identify and possibly acquire data sets that the EW team is not at liberty to share. In this case, the EW team noted two necessary steps: a review of data sharing agreements, particularly for household survey data, as well as an evaluation of the LOE to clean and upload the data sets into the FDW.

FDW/FDE functionalities. The EW team proposed the following changes to data extraction processes:

- ***Restructure data.*** Certain FDW data series (i.e., IPC, population, crop, and nutrition data) would be more useful if restructured, thus eliminating the need for the EW team to further manipulate the data once extracted.
- ***Multiple data series per geography.*** To eliminate the need for the EW team to stitch together data from several data domains to focus on a country or region, the FDW should allow users to select a geography and bring together all desired data within that space. Particular series often jointly analyzed include livelihood zones and prices with commodity balances.
- ***Automate time-consuming processes.*** Candidates for automation through the FDW included: data ingestion and sharing; the data entry components for price projection templates; scripts for household data analysis; visualizations on the website; and production of the shapefiles on the FEWS NET Data Center.
- ***Refreshable spreadsheets.*** The EW team requested a refreshable spreadsheet for informal cross border trade and commodity balances, which would allow analysts to link data series with price trend and commodity availability visualizations. Additionally, the EW team requested a specific fix to the existing FDW refreshable prices spreadsheet that currently creates metadata problems.
- ***FEWS NET website.*** The EW team expressed interest in automating monthly monitoring reporting, as long as they could continue to curate content and complement it with appropriate contextual information. An example provided was online tools that allow users to compare commodity balances from the current year against the El Nino year. The EW team cautioned the automation of certain reporting, including price projections, because the EW team does not fully base projections on mathematical models—unlike WFP and FAO—and rely on expert knowledge of local contexts.

[35] Note that some of these indicators are beyond the Hub's scope of work for Y2, as indicated in *Volume 1.*

Research frontiers. The EW team proposed the Hub identify new approaches for processing spatial data sets (e.g., crowd "tagging" of images or using machine learning to estimate total area planted).

FEWS NET MANAGEMENT TEAM PILLAR 3 CONSULTATION

On April 14, 2020, the Hub interviewed ***Phil Steffen,*** the Program Manager for FEWS NET's Pillar 3 (P3) *Analysis of the Dynamics of Food, Nutrition and Livelihoods Security*. The Hub's COR, Romaine Williams, and Program Manager, Kevin Coffey, also participated in the interview. Discussion covered questions related to what Steffen, as well as future P3 implementing partners, might expect from both the Hub-run FEWS NET website and the DMP.

Steffen described P3 as a vehicle to support various USAID clients—including FEWS NET, the BHA, or any operating unit within USAID—that may issue TO(s) to contractors in an effort to understand and/or address the complex causes of food insecurity and vulnerability. A fuller list of illustrative topics is available in the Request-for-Proposals, but some examples of P3 work include:

- Updating FEWS NET's foundational background reports and assessments;
- Responding to requests to address data and knowledge gaps identified by other USAID offices, by synthesizing state-of-art knowledge rather than conducting original analyses; and
- Developing new methodologies to analyze the drivers of food insecurity, including do-it-yourself analytical tools.

Relationship to Hub Platforms. P3 contractors should be able to use the DMP and FEWS NET website to access the tools and information required for these tasks. Inputs to P3 analyses likely include data sets, literature reviews, focused interviews, and webinar participation. On a case-by-case basis, the Hub may directly support P3 studies.

Likewise, P3's outputs—such as data, tools, and reports—should be available, as appropriate, on Hub platforms. P3 audiences include the FEWS NET EW team, other USAID operating units, and the global food insecurity field. Specific data and indicators likely needed by P3 analysts and potential candidates for ingestion into the FDW include:

- Crop production data, currently accessed from FAOSTAT;
- Population and demographic data, currently retrieved from UN's online holdings that allow parameterized projections;
- Migration data;
- Urban food security indicators; and
- Factors threatening food security (e.g., state fragility, conflicts, and pandemics).

P3 is also likely to require data from well-known, reputable, international organizations, such as the FAO, WFP, Center for Strategic and International Studies, and National Rescue Committee, but whose data may not be candidates for inclusion in the FDW.

P3 analysts likely require FDW functionalities that allow tabular and visual presentations across time, and in particular, across previous reference periods (e.g., other famines or pandemics). P3 analysts

may also need to track the evolution of methodologies over time, including the development of different indices and scales.

Longer term vision. Steffen reported that he was familiar with the FEWS NET website, but not yet with the DMP. He noted that at present, the website is primarily a repository for narrative reports on FEWS NET presence countries. He stated that if FEWS NET is to truly be the "gold standard" for early warning analysis, it must prioritize making data an international public good. Furthermore, the DMP must be robust, capable of handling a significant amount of data, and maintain tools that allow analysts external to the project to develop their own food insecurity scenarios. The DMP's information must be unbiased and useful to inform the types of decisions that lead to better food security outcomes. He reported that a combination of the data in the FDW and such analytical tools could prove to be powerful at impacting policy making.

BUREAU FOR HUMANITARIAN ASSISTANCE

DDATA CONSULTATION

On March 19, 2020, the Hub met with ***Langdon Greenhalgh***, Project Manager, and ***Mark Slezak***, Personal Services Contractors Generalist, both of whom are in USAID's DDATA unit. In the aftermath of the Ebola outbreak, USAID created DDATA to create a stronger evidence base for disaster management at USAID, which includes how data are leveraged to produce analyses and reporting. At present, DDATA focuses on supporting ad hoc analyses and reviews within USAID, attaining higher quality data to inform humanitarian response, and enhance partner capacity.

Relationship to Hub platforms. Both informants reported using the FEWS NET Data Portal on the fews.net website to review maps, IPC shapefiles, and price data. They indicated that the background narrative in the PDF reports is valuable to provide context to the visualizations, but stated that more sub-national data are always needed. However, Venezuela, which is their current country of focus, is not a FEWS NET presence country, so they are relying on other data sources. They are both interested in exploring the FDW to identify additional uses of FEWS NET data for a broader analytical base to inform future DDATA planning.

Data sources and software applications. Greenhalgh and Slezak stated they primarily use Tableau and Power BI to visualize and link boundary data to geospatial data sets, including climatological data (e.g., earthquakes and floods) and hazard data (e.g., risks related to infrastructure and disease outbreak), collected from sources such as the Humanitarian Data Exchange, the REACH initiative of the Assessment Capacities Project, the displacement tracking matrix from the International Organization, and data portals from Humanitarian Coordination clusters. Because most of these data are managed on an ad hoc basis through spreadsheets, and less frequently as narrative in PDFs from on-the-ground partners, they expressed interest in accessing structured and tagged data in a central data platform such as the FDW .

Current needs for support. Greenhalgh and Slezak indicated the Hub could play a role in supporting data and documentation management for sectors that drive or indirectly impact food insecurity (e.g., creating metadata schemas to handle ambiguous, high level, or non-tabular crop production data).

They identified the collective challenge of meeting data management needs from technical, project, and management perspectives. Across USAID, teams are continuing to identify the use case for certain data, but organization overall would likely benefit from stronger technical guidance and better data architecture.

USAID BUREAU FOR RESILIENCE AND FOOD SECURITY

ANALYTICS AND LEARNING DIVISION

On March 13, 2020, the Hub met with ***Baboyma Kagniniwa,*** Senior Data Scientist in the Analytics and Learning Division of the BRFS. Kagniniwa uses data science methods to demonstrate the impact of BRFS interventions to "tell the story" of USAID's programming, and to identify potential beneficiaries of future programming. He analyzes what's going on in individual countries, compiling information across the different donors. He provides data to a number of colleagues, including:

- Senior staff who request key, high-level statistics at the global or regional level to testify before Congress or attend global events;
- Communications specialists telling the story of how BRFS impacts on-the-ground situations (e.g., how agricultural investments impact prices, or how IPC levels change following programming interventions);
- Senior monitoring and evaluation specialists who report on local contexts and thus require IPC levels; and
- Analysts like himself who request raw data.

Relationship to Hub Platforms. Until recently, Kagniniwa accessed FEWS NET information either through the website or upon special request for data. After receiving direct access to FDW, he indicated his experience with FEWS NET information substantially improved. Kagniniwa typically turns to FEWS NET sources when he seeks to demonstrate the impact of USAID programs, particularly in drought-affected areas. He regularly uses IPC data, often at the livelihood zone or other disaggregated level; agriculture data, including natural food stocks, prices, and trade data; and biophysical data such as the Normalized Difference Vegetation Index (NDVI) from the USGS website.

Suggested improvements to FDW. Kagniniwa stated that the location of data and associated resources (e.g., how data are stored and organized, as well as its distribution channels) are important to his work. He provided the following recommendations to improve the FDW:

- ***Expand nutrition data.*** Though he reported frequent use of nutrition data that are unavailable in the FDW (e.g., USGS stunting indicators), he would like to see a self-service dashboard on FDW, harmonized and linked to other frequent FEWS NET reporting. With the reorganization, USAID has created a Nutrition Data Center which might be a useful partnership for FEWS NET.
- ***Feature downloadable maps.*** He is often tasked with reverse-engineering FEWS NET maps, which requires significant effort even when using the FDW. He suggests making all the maps FEWS NET produces, including regional, national, and sub-national analyses, available separately from the report PDFs to allow further exploration of the data.

- ***Facilitate data extraction.*** Kagniniwa suggests disaggregating data from FEWS NET reporting, letting users define data extraction timeframes and, prioritize implementing the Application Programming Interface to facilitate access to a larger quantity of data.
- ***Link to research and use cases.*** Lastly, he advised the Hub to identify USAID's burning research and then work to make the data fit those use cases. An example would be monitoring and evaluating USAID's investments to outcomes and impacts.

FOOD SECURITY MONITORING AND EVALUATION GROUP, ANALYTICS AND LEARNING DIVISION

On March 17, 2020, the Hub met with ***Kiersten B. Johnson, PhD,*** Senior Researcher at USAID BRFS. Johnson regularly uses both social survey and remote sensing data sources to better understand development context "on-the-ground". She regularly responds to ad hoc data and analytical requests, and played a key role in expanding the use of remote sensing and location-based data within BRFS.

Relationship to Hub platforms. Johnson has visited fews.net in search of data and information but finds navigation difficult, and country coverage differs from what she needs for her BRFS analyses. She is not familiar with the FDW.

Data sources. Johnson noted that with the reorganization to BRFS, the former Monitoring, Evaluation, and Learning Division would become the Analytics and Learning Division. With this shift, she anticipates a focus from historically simple metrics to a more nuanced understanding of the dynamics driving the outcomes they seek. This shift will open up opportunities for BRFS to collaborate more closely with FEWS NET, which has rich geospatial agroclimatic and socio-economic data sets.

She is particularly interested in tracking and understanding the environmental and natural resource factors which influence agriculture. When BRFS tracks the impacts of project interventions on crop yields, it needs to take into account the other confounding factors such as agroclimatic events by incorporating data on El Nino events, rainfall and seasonal variability (using the Standardized Precipitation Index and the Normalized Difference Vegetation Index).

Likewise, she can see applications for incorporating spatial information on prices, trade and livelihoods into their assessments of impacts both on farmers and traders. FEWS NET's Livelihoods data, which she has yet to use, would be of interest as BRFS selects and implements in its "Zones of Influence."

Software considerations. At present, analysts in her office are doing simple comparisons of data. They lack the models and the software to conduct more sophisticated multivariate analyses needed to disentangle the dynamics described above. She has been advocating to use data, such as available through FEWS NET, and to expand analytical methods (perhaps in R or Python).

Collaboration opportunities. Johnson has appreciated the opportunity to learn about the utility of remotely-sensed data from the FEWS NET Management team. She has also appreciated working with the USAID SERVIR project's working group focused on NASA's Earth Observing System.[36] She

[36] https://servirglobal.net/

welcomes an opportunity to be active in such communities of practice. As her office is short-handed, she would appreciate interacting with other experts who can judge the quality of data.

In addition, she would appreciate a FEWS NET presentation on the types of data the project can make available to others at USAID. An FDW data presentation would be of particular interest.

BRFS RESILIENCE SPECIALISTS

On March 18, 2020, the Hub met with two resilience specialists:

- ***Andre Mershon,*** Resilience Adviser and Country Support at USAID BRFS, who provides support to resilience countries, also FEWS NET presence countries, in West and East Africa; and
- ***Jami Montgomery,*** Division Chief, Center for Resilience, at USAID BRFS who provides community assistance to inform USAID programming decision-making.

Both interviewees also indicated that their work is increasingly examining the impact of micro shocks (e.g., illness) above and beyond the larger macro shocks (e.g., climate and markets) that can undermine household resilience.

Relationship to Hub Platforms. Both resilience specialists reported extensively using the FEWS NET Data Center on the fews.net website, as well as drawing from the PDF reports. Mershon recently used FEWS NET information, including IPC levels, to design a survey to understand household resilience in the aftermath of recent shocks. Though Montgomery previously accessed FEWS NET data on climate trends, including NDVI, through email notifications and livelihood zone map downloads, she now prefers to submit requests to be sent data and information.

Humanitarian responses data in the FDW. To understand how resources are allocated, they asked Gary Eilerts, who was at that time the FFP FEWS NET Surge Officer, to compile data on food aid as a proxy for humanitarian assistance at the country, and where possible, sub-national levels. This detailed data helped several offices at USAID—including Office for Coordination of Humanitarian Affairs, Emergencies and Natural Disasters, and others—to prioritize programming in countries where FEWS NET repeatedly identifies acute food needs.

That long term cumulative data is still being used for a variety of other analyses. For example, BRFS is working with TANGO International on understanding the returns to development investments in terms of the amount of humanitarian investment avoided.

Additionally, Montgomery remains interested in developing tools to better understand and visualize their data holdings, recognizing the potential value in developing longer-term supportive relationships between FEWS NET and other USAID Offices.

Suggested improvements to Hub platforms.

- ***Time series maps.*** Montgomery expressed interest in seeing subnational agroclimatology maps and seasonal forecasting that include trends over time (e.g., drought and flooding frequencies and/or temperature changes over the last decade). This type of data and

associated visualizations would help facilitate USAID planning of humanitarian responses over the long term (e.g., five-year programming).

- ***Simplified guide to agroclimatology information.*** Because Montgomery primarily draws from the narrative analyses found in the PDF reports, rather than the data itself, she found the language in agroclimatology reports to be too complex. She recommended revising the reporting language to be easily understood by non-technical audiences.
- ***Open source data.*** She stated that data from non-permission environments is critical to their analyses because USAID missions vary in their mapping skills and software.