Internet Computer Protocol

CALL FOR ACTION

Simple Web Server ICP Project

Architecture, Insights, Outcomes and Recommendations

Call for action

TOC

Call for action

Introduction

This project aims to build a web server that collects financial data from Yahoo Finance and displays it for users. Additionally, the web server provides a simple subscription service that allows users to subscribe via email for updates or alerts. The focus is on integrating data collection with user notification capabilities in a streamlined way.

Web Server ICP Project: Architecture and Insights

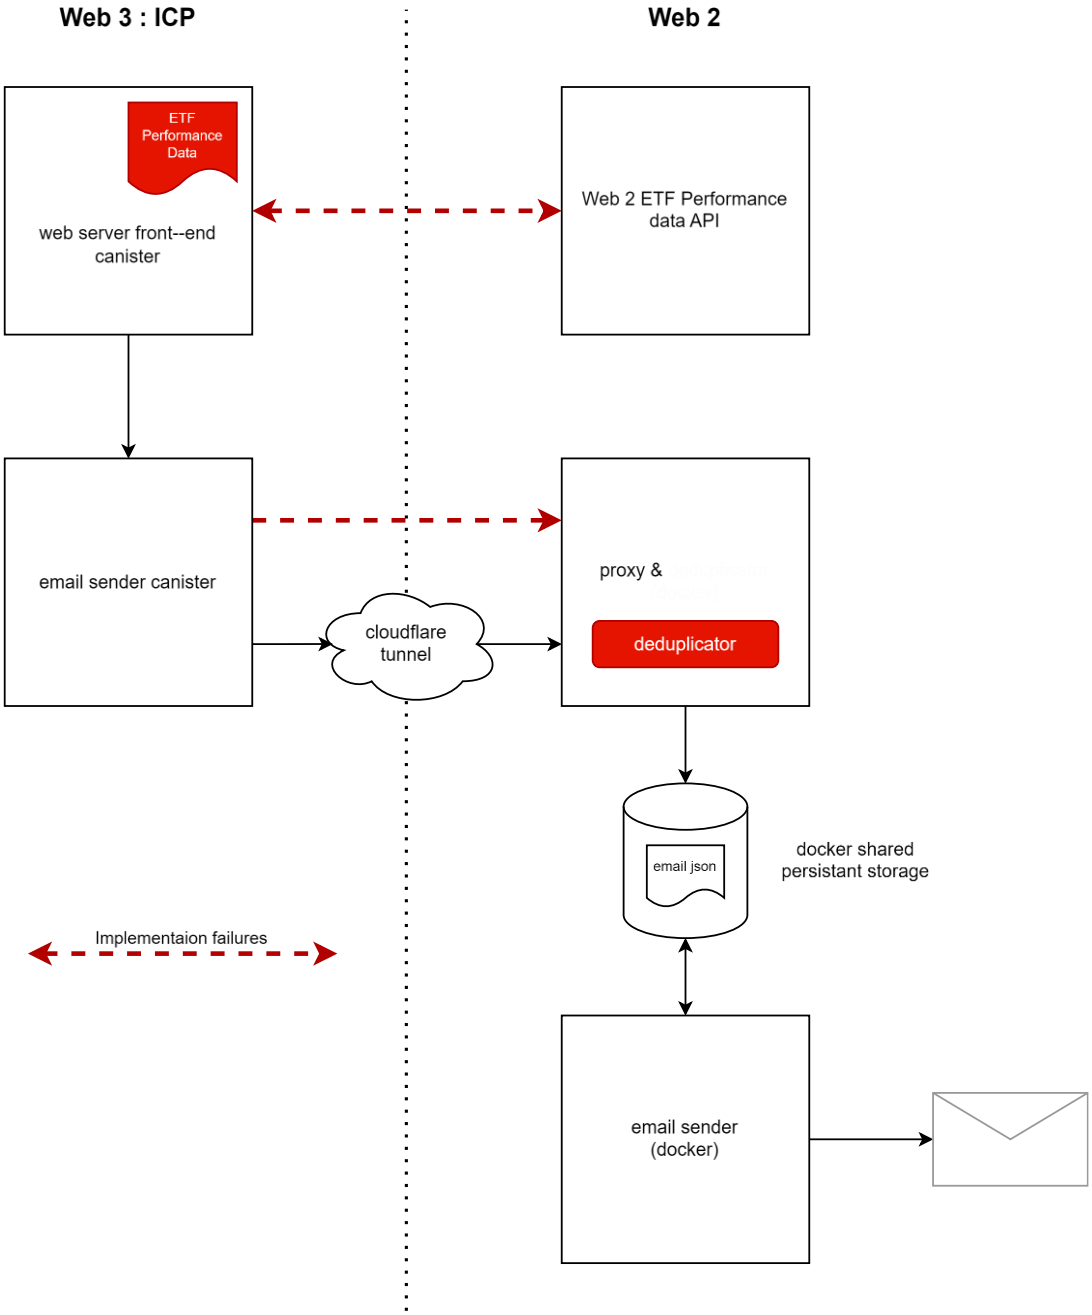

Call for action

Full Data Flow Description

ICP Layer

- The Web Server Canister contains hardcoded ETF performance data embedded during canister deployment or upgrade. This design circumvents ICP's limitations in fetching dynamic external API data, which causes consensus failures due to inconsistent responses across replicas.
- An external off-chain service runs periodically outside ICP to fetch ETF data from Yahoo Finance. This data is transformed into a hardcoded initial dataset baked into the Web Server Canister, ensuring consistent display of ETF performance graphs.
- Refreshing the data requires rerunning the external fetcher and redeploying or upgrading the canister with updated static data.
- The Email Sender Canister handles email notifications, triggered by the Web Server Canister as needed.

Email Notification Flow

- The Email Sender Canister sends emails through a Cloudflare Tunnel to your on-premise infrastructure. This tunnel is necessary because ICP cannot directly connect to on-premise servers, requiring this secure, outbound-only connection.
- Setting up and maintaining the Cloudflare Tunnel is an additional operational overhead and a potential point of failure.
- On-premise, a Docker container acts as a proxy and deduplicator, receiving email JSON files from multiple email sender replicas (13 in your case). It consolidates these into a single canonical JSON file within persistent storage.
- A second Docker container polls the buffer every 60 seconds and sends the email to your Gmail account using app authentication, ensuring reliable, authenticated delivery.

Architectural Constraint, Workaround, and Overhead

- ICP is not suited for collecting real-time API data due to consensus constraints; static data and canister rebuild are currently required to update displayed information.
- Direct on-premise server access requires the Cloudflare Tunnel workaround, adding setup and maintenance complexity.
- Email JSON files sent from the Email Sender Canister must be deduplicated via the on-premise proxy to manage multiple ICP replicas sending the same data.

Call for action

Outcomes and Recommendations

- Enhance ICP's integration with Web 2.0 API components, focusing on consensus/determinism challenges so canisters can consume external data reliably without complex workarounds.
- Improve native integration with Web 2.0 systems to eliminate the need for proxy layers that handle deduplication and add operational overhead.
- Facilitate seamless integration with on-premise infrastructure to avoid reliance on components like Cloudflare Tunnel that introduce complexity and single points of failure.
- Addressing these issues is essential for ICP to reach massive adoption and enterprise readiness.
- The Utopia project, while promising, segments ICP into two environments (public and private), which may double development effort. A more desirable approach is a unified platform that transparently operates across both public internet and private on-premise networks.

Call for action

Acknowledgment and Call for Collaboration

I am at the early stages of exploring and developing on the Internet Computer Protocol (ICP). While this project represents my initial efforts to understand its capabilities and limitations, I am eager to deepen my knowledge and improve my implementations.

I would greatly appreciate any insights, experiences, or solutions from fellow developers and experts who have faced similar challenges or have built robust architectures on ICP. Sharing knowledge and collaborating will help overcome current obstacles and contribute to the growth of the ICP ecosystem.

If you are willing to share your expertise or discuss potential approaches to the problems encountered in this project, please feel free to reach out.

Jan Affolter,
Digital Solution Designer (AI & Analytics),
11 November 2025

LinkedIn: https://www.linkedin.com/in/janaffolter/