TigerGraph

Customer Success Story

Xandr

FOUNDED IN: 2018
LOCATION: NEW YORK
EMPLOYEES: 1,800+
AREA OF FOCUS: ADVERTISING

Xandr Builds an Identity Graph for Millions of Customers for Entity Resolution using TigerGraph

Introduction

Xandr is the advertising and analytics division of AT&T's WarnerMedia. The company operates an online platform called Community that connects advertisers, publishers, and consumer media brands. Community is built on Xandr's powerful technology and consumer insights, enabling buyers and sellers to drive better business outcomes. Xandr was founded in 2018, employs approximately 1,800 employees, and is headquartered in New York City.

"The fragmentation of the user journey necessitates the need for graph database technology where I can retrieve multiple entries without having to have an endless number of column joins."

Dr. Abraham Greenstein, Sr. Director, Data Science, AT&T

The Challenge

Xandr wanted to create a digital advertising platform, Community, that would include over 15 WarnerMedia properties such as Cinemax, CNN, HBO, TNT, and more. However, since each of these brands has millions of consumers, disambiguating data in the proposed new platform would be a challenge. Not only that, Xandr wanted to combine all of this information in a way that would enable a better platform for advertisers, one capable of creating personalized commercials, with different promotions being delivered to individual viewers simultaneously. This would be an industry-first.

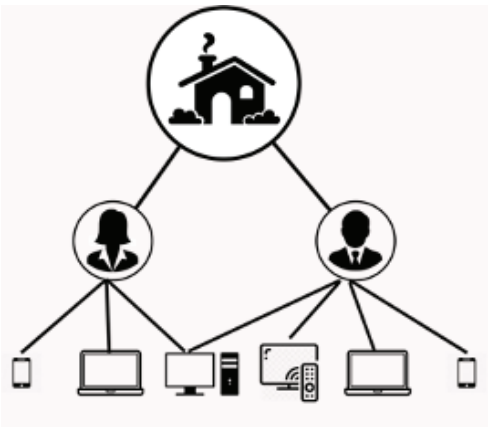

An Identity Graph stitches together different identifiers into a unified view of people, the households they belong to and devices they use.

The Solution

TigerGraph is enabling Xandr to merge data from silos across the WarnerMedia universe and build the first, and largest, identity graph of its kind in the advertising industry. TigerGraph's ability to integrate datasets and resolve ambiguous entities at scale is allowing Xandr to uncover insights hidden in its data. Specifically, Xandr is able to identify attributes associated with people, devices, households that provide marketers the ability to target audiences with customized commercials aligned to their interests and needs, resulting in better advertising performance.

"We have a large distributed graph with over 5 billion vertices (entities) and 7 billion edges (relationships). Every single day we make up to 1 billion updates. Every time we run our Identity Resolution Algorithm we create 300 million more vertices and a billion more edges. We need to be able to scale horizontally and TigerGraph allows us to do that."

Chinmay Nerurkar, SSE2 Team Lead (Associate Director), Data Science Engineering, Xandr

The Results

Xandr's SSE2 Team Lead (Associate Director), Data Science Engineering, Chinmay Nerurkar, says that identity graphs are important to advertisers as they stitch together different identifiers into a unified view of people, households and devices to enable cross-device and converged addressable advertising. To this end, Xandr has developed an identity graph with over five billion vertices and seven billion edges implemented with a ten node distributed cluster, each node with 400 GB RAM and 48 cores using TigerGraph. This is enabling Xandr to increase revenues by attracting companies to advertise on its platform, instead of using competitive products or other outmoded approaches.

ENTITY RESOLUTION USING CENTRALITY ALGORITHMS. CENTRALITY IS A MEASURE OF A VERTEX'S IMPORTANCE.

- **Distributed graph with 5+ billion vertices and 7+ billion edges**
- **Up to 1 billion daily graph updates from input**
- **300 million vertices and 1+ billion edges created by the algorithms**
- **Built a 10 node TigerGraph cluster. Each node has 48 cores, 400GB RAM, 3GBps NVMe storage**
- **Running BFS-style algorithms, like Label Persistence, spanning over a large distributed graph is extremely memory intensive**

Source: "Identity Resolution at Scale" by Chinmay Narurkar and colleagues, presented at Graph + AI World 2020 conference. Watch the session: https://info.tigergraph.com/graph-ai-world-xandr

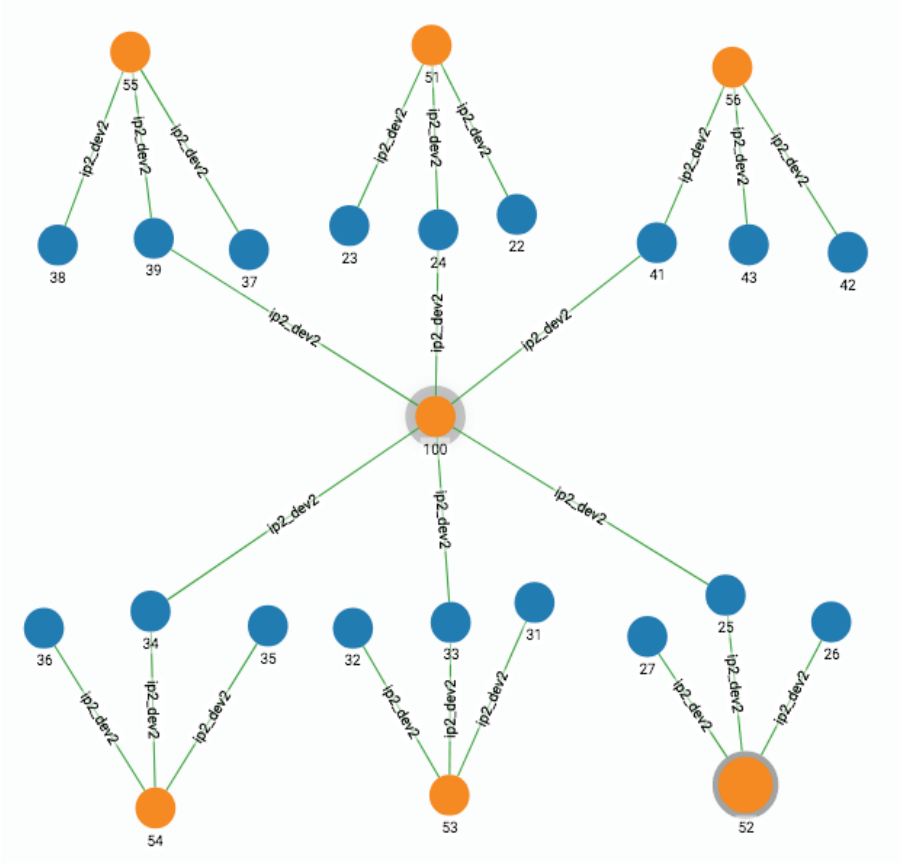

"You need to exercise graph thinking. If you try to build your graph with a relational model in mind that relies heavily on complex joins and where clauses, you're not going to get very far. TigerGraph stores related vertices adjacent to each other and allows you to traverse these relationships extremely fast"
Chinmay Nerurkar, SSE2 Team Lead (Associate Director), Data Science Engineering, Xandr

"Our graph gives us cleaner, better data to feed into our learning algorithms. We start looking at the household level. Not how many times a user has seen an ad on a given device, but how many times they've seen an ad through all their devices. That cleaner data makes the ML(machine learning) work better"
Dr. Abraham Greenstein, Sr. Director, Data Science, AT&T

Some of Our Customers

AMGEN · intuit · wish · 国家电网公司 STATE GRID CORPORATION OF CHINA · opencorporates · Pagantis · IPPEN DIGITAL PLATTFORM · Kickdynamic

- Real-time fraud detection at 7 out of the world's top 10 global banks
- Care path recommendations for 50 million patients
- Personalized offers for 300 million consumers
- Energy infrastructure optimization for 1 billion people

"Once we have everything built in graph, we can react to changes in real-time. Graph is at the centre of everything we do."

Dr. Jay Yu
Distinguished Engineer and Architect,
Intuit

"With TigerGraph we can join sources of data together and make connections within the data that previously we couldn't. We can now answer questions that, for the last 20 years, we didn't think were possible to ask."

Harry Powell, Director of Data & Analytics,
Jaguar Land Rover

About TigerGraph
TigerGraph is the only scalable graph database for the enterprise. TigerGraph's proven technology connects data silos for deeper, wider and operational analytics at scale. Seven out of the top ten global banks use TigerGraph for real-time fraud detection. Over 50 million patients receive care path recommendations to assist them on their wellness journey. 300 million consumers receive personalized offers with recommendation engines powered by TigerGraph. The energy infrastructure for 1 billion people is optimized by TigerGraph for reducing power outages. TigerGraph's proven technology supports applications such as fraud detection, customer 360, MDM, IoT, AI, and machine learning.

For more information visit www.tigergraph.com and follow us at: Facebook Twitter LinkedIn

Contact us at sales@tigergraph.com

TigerGraph
3 Twin Dolphin Drive, Suite 225
Redwood City, California 94065

Get Started for Free at Tigergraph.com/Cloud

TigerGraph Cloud graph database as a service is built for agile teams who'd rather be building innovative applications to deliver new insights than managing databases.

Cloud Starter Kits

TigerGraph Cloud Starter Kits are built with sample graph data schema, dataset, and queries focused on specific use cases such as fraud detection, recommendation engine, supply chain analysis and/ or a specific industry such as healthcare, pharmaceutical or financial services.

STARTER KIT	OVERVIEW
COVID-19 ANALYSIS	Detect hubs of infection and track the movements of potential spreaders
CUSTOMER 360-ATTRIBUTION & ENGAGEMENT GRAPH	Create a real-time 360 view of the customer journey for attribution and engagement insights.
CYBERSECURITY THREAT DETECTION-IT	Block cybersecurity threats by detecting interconnected events, devices and people
ENTERPRISE KNOWLEDGE GRAPH (CORPORATE DATA)	Analysis of corporate data including investors and key stakeholders.
ENTERPRISE KNOWLEDGE GRAPH (CRUNCHBASE)	Knowledge graph examples featuring crunchbase data with startups, founders and companies.
ENTITY RESOLUTION (MDM)	Identify, link and merge entities such as customers with analysis of attributes and relationships.
FRAUD & MONEY LAUNDERING DETECTION	Multiple types of fraud and money laundering patterns.
GSQL 101	Introduction to TigerGraphs powerful graph query language.
HEALTHCARE GRAPH (DRUG INTERACTION/ FAERS)	Healthcare example focused on public (FAERS) and private data for pharmaceutical drugs.
HEALTHCARE–REFERRAL NETWORKS, HUB (PAGERANK) & COMMUNITY DETECTION	Analyze member claims to establish referral networks, identify most influential prescriber's and discover the connected prescriber communities.
MACHINE LEARNING & REAL-TIME FRAUD DETECTION	Mobile industry example for detecting fraud in real-time and generating graph-based features for training the machine learning solution.
NETWORK & IT RESOURCE OPTIMIZATION	Network and IT resource graph for modeling and analyzing the impact of the hardware outage on workloads.
RECOMMENDATION ENGINE (MOVIE RECOMMENDATION)	Graph-based movie recommendation engine built with public data.
SOCIAL NETWORK ANALYSIS	Social network example for understanding and analyzing relationships.
SUPPLY CHAIN ANALYSIS	Example covering inventory and impact analysis.