

In Seattle, behind-the-scenes Facebook team wrangles digital deluge in massive cloud

August 23 2017, by Matt Day, The Seattle Times

More than 2 billion people log into Facebook every month. Every day, the social-media crowd uploads billions of photos, calls up hundreds of millions of hours of video, and fires off a prodigious flurry of likes and comments. Somebody has to store that slice of humanity's digital record.

Much of that task falls to Surendra Verma, a Seattle engineer who for more than 20 years has been building software that files away and retrieves large volumes of data.

Verma leads the storage team at Facebook, the group charged with making sure that the social network can accommodate its daily deluge without losing someone's wedding photos.

Most of that unit is based in Seattle, part of a workforce that today numbers 1,600 people, up from just 400 three years ago. That makes Facebook one of the fastest-growing technology companies - outside of Amazon, anyway - in the city.

While Facebook employees work on a wide range of products in Seattle, the office has developed a specialty in the geeky realm of systems software.

About a quarter of the Facebook engineers in Seattle work on the company's infrastructure projects, the tools to transmit, store and analyze the growing heap of data people feed into the social network.

That's a common trade in the region, where Amazon Web Services, Microsoft and Google are all building their own clouds - giant, globe-straddling networks of [data centers](#) and the software that manages them.

Facebook could have built its products on computing power rented from those cloud giants, but it decided to build its own tools, from custom hardware designs all the way to mobile applications. Supporting Facebook's network are nine massive data centers - a 10th, in Ohio, was announced earlier this month.

Facebook's cloud is different from the others' in that it's designed to support just one customer: Facebook's own apps.

They happen to be some of the most widely used pieces of software in the world - and their use keeps expanding.

Verma is an Indian-born engineer who got his start at IBM before moving to Microsoft, where he worked most recently on Windows file systems. He joined Facebook in 2015.

By then, the Seattle office had come to pilot Facebook's storage-software efforts. "We could find some very good engineers here," he said. "And so a bunch of these projects started, and just got momentum from there."

His team's job, he said, is to be invisible to Facebook's product groups, letting the company's other engineers build whatever they can think up on top of a reliable foundation.

But a wave of new services, and the rapid growth in the number and quality of videos and photos that people share, is putting a huge burden on the network's infrastructure.

Facebook in January 2016 started a monthslong rollout of live video streaming, a milestone in the social network's effort to compete with popular streaming services.

Then Instagram Stories, aimed at competing with the ephemeral photo-sharing application Snapchat, launched last August, and quickly made its way to 250 million monthly users. (Facebook bought Instagram in 2012.)

Up next: live video broadcasts on Instagram, called Live Stories, a feature the product group hoped to launch before Thanksgiving.

Verma's team wasn't sure its systems could meet that deadline, and negotiated a few days of delay on the planned start date. After scrambling, the team scraped together the internal storage space needed to accommodate the new feature, which went live Nov. 21.

"We looked very hard at our capacity, everywhere," Verma said. "We scrounged everything we could, looked at nooks and crannies."

One of the people doing the looking was J.R. Tipton, a manager on Verma's team.

Tipton left the Chicago suburbs and came to Seattle in 2001 for the same reason as thousands of others in the 1990s and 2000s: a job offer from Microsoft.

Fifteen years later, he became part of another phenomenon reshaping the region, opting to leave maturing Microsoft for a job with a younger tech giant in the area.

"I wanted an adventure," Tipton said of his move to Facebook last year.

Tipton and Verma are among the 640 people at Facebook in Seattle who,

on LinkedIn, list past experience at Microsoft.

Seattle, home to Boeing's legions of engineers long before Microsoft came along, has never really been a one-company town. But the range of options for technologists has ballooned in the last decade, with Amazon uncorking another Microsoft-like growth spurt, and Silicon Valley giants seeding local outposts to scoop up some of the talented software developers here.

When Facebook set up shop in Seattle in 2010, it was the company's first U.S. engineering office outside its Menlo Park, Calif., headquarters. Last year, Facebook Seattle, which then numbered about 1,000, moved into nine floors of custom-designed office space. That was followed in short order by two more leases that would give the company enough space for about 5,000 employees in Seattle.

Today, Tipton works on the Facebook system that is the last line of defense keeping people's photos and videos from getting wiped out by a software error or power outage at a data center.

That would be cold storage, the backup system storing redundant copies of all of the photos, videos and other data Facebook stores.

The engineers who designed the system in 2013 wanted to build something more efficient than the typical rack of servers, which suck up energy both to keep their hard drives powered on and to run the fans and air-circulation systems that keep the stack of electronics from overheating.

The company landed on a design that leaves most hard drives in Facebook's rows of servers powered down until needed, hence the "cold" in cold storage.

The software built to run that system was designed to handle plenty of growth in the amount of data people were throwing at Facebook.

Each cold-storage room in Facebook's data centers was built to house up to one exabyte of data. That's 1,000 petabytes, or 1 billion gigabytes, or the combined hard drive space of 500,000 top-of-the-line MacBook Pro laptops.

"At the time, it seemed laughably huge," Tipton said of the software layer built to manage all of that.

Last year, it became clear it wouldn't be enough, the result of growth in both the number of Facebook users and the launch of data-intensive features like Instagram's Live Stories.

Tipton and his Seattle-based group spent the first half of 2016 taking a long look at the cold-storage software, which was starting to show its age in a world of 360-degree video. They could invest more time and energy in keeping the old architecture working, or rebuild.

They opted to rebuild.

"It has to be much, much bigger," he said.

Tipton has metaphors at the ready to describe that continuing challenge, the fruit of years explaining his day job to nontechies.

"You stand on top of a hill and see the buffalo herd coming," he said.

And if that doesn't sound harrowing enough, he tries again: "We're basically laying the railroad track down as the train is coming."

©2017 The Seattle Times

Distributed by Tribune Content Agency, LLC.

Citation: In Seattle, behind-the-scenes Facebook team wrangles digital deluge in massive cloud (2017, August 23) retrieved 10 April 2024 from <https://phys.org/news/2017-08-seattle-behind-the-scenes-facebook-team-wrangles.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.