Contradictions in advanced cloud-computing research

September 9 2013

On first appearances, it might seem as if researchers at the University of California, San Diego Center for Networked Systems are working at cross-purposes to one another.

Many CNS researchers are exploiting the rise of mobile and cloud computing to make data available anytime, anywhere and to anyone, at faster speeds and greater reliability. Meanwhile, their colleagues across the hall (and sometimes even across the desk) strive to make that data impossible for anyone to access at any time—or at least anyone who is not an authorized user.

That philosophical tension is part and parcel of what CNS researchers call "the beauty of the decentralized system," which has grown to dominate the way people create, share and store information. At its two-day, twice-yearly research review—held earlier this month on the UC San Diego campus—CNS demonstrated its role as a major player in the ongoing process of designing, managing and improving data center and wide-area networks. CNS graduate students were a significant presence at the review as well, both as presenters and as participants in a poster session and reception that showcased their work.

"UCSD is just an awesome place," enthused research scientist George Porter, who is also the associate director of CNS. He said that the principal investigators at CNS—many of whom are based at the Qualcomm Institute—"are known for making big contributions in designing scalable, fault-tolerable networks and understanding how

networks work. They're known for developing and designing next generation storage technologies and they're also making advances in terms of security writ large, from understanding how spam works to understanding how to make the computerized systems in cars safer and more reliable."

If there's any occupational hazard to working in the field of networked systems, said Porter, it's that researchers sometimes get so entrenched in their own piece of the puzzle that they don't always stop to think about what the whole puzzle should look like.

"The research review is both a way for us to advertise and expand our influence and impact and also as researchers to connect with our member companies and others in industry to understand their problems," he said. "We don't always spend a lot of time talking about those big abstract topics when we're working from day-to-day."

One of those member companies is Google, which recently recruited CNS researcher and Computer Science and Engineering Professor Amin Vahdat to help re-architect one of its Wide-Access Networks (WAN), known as B4. The B4 network connects Google's many data centers, and a second Google WAN makes it possible for the public to conduct Internet searches, download YouTube videos, etc.

Vahdat presented a summary of Google's efforts to economize and improve B4 as one of the keynote talks at the research review, which was attended by several representatives from CNS' half dozen other industry partners (Cisco, Microsoft and Oracle among them).

One of the challenges Vahdat's team faced in redesigning the B4 network stems from the sheer amount of data being shared via the WAN at any given time. According toWired magazine, Google serves a full 25 percent of Internet traffic in North America. Because users want that

traffic to flow as quickly as possible and do not want to be exposed to the effects of individual link failures, Google's WAN must consequently 'overprovision' data paths, or take more bandwidth than it really needs. Vahdat explained that for a speed of 10 Gb/sec data, for example, the path must be capable of handling 30 to 40 Gb/sec.

Despite the high demand on its networks, existing WANs cannot run at 100 percent capacity because doing so can cause packet failures (and subsequent data loss). When a path fails, noted Vahdat, "we don't know which path to provision" because the system is so decentralized.

"A human has to remember a mental model of how a path that goes down will impact the rest of the network," said Vahdat. In an effort to devise a work-around, "we asked ourselves: 'Can we differentiate between high-priority, no-loss traffic and traffic that can experience some loss for an hour but later make up the difference?'"

The traditional, decentralized approach to networking takes an "all bits are created equal" view, which leads to a situation where "apps that could benefit from additional bandwidth are denied so other bits don't get dropped," he explained. "But on virtually any network, the bulk of bytes aren't high priority."

A researcher who is trying to upload large amounts of data to the cloud over time, for example, might not need that data transferred with the same per-packet requirements as someone live-streaming a YouTube video. What is certain, he said, is that "Internet service providers are going to have a tough time determining prioritization, so we need new service models."

He added that such models will require an overhaul of the existing software/hardware WAN architectures, which are impractical for delivering the necessary bandwidth on a global scale. Software-defined

networking solutions like B4 are one approach to determining prioritization, and would likely reduce costs for ISPs because prioritization of data would allow the WAN to be run more cheaply and would require less overprovisioning.

Given all of the challenges associated with handling this data avalanche, it might seem that storing data in the cloud—rather than in a computer under one's desk—is more trouble than it's worth. And it's true that when data and computing are opened up to the Internet at large, bad things can happen. Online services fail ("and there go all your wedding photos," said Porter). But hard drives fail, too, and typically more often. The benefits of cloud computing are so substantial, in fact, that Porter predicted the trend "will only become deeper and more broad" in the coming decade.

"The ability to replicate your data across the organizations in the cloud gives you a lot more reliability than you yourself have if you were to store your data only on a USB drive," he said. "Not only do you get more computing power from the cloud, your data now becomes more interesting because one can view it within a social context or based on conditions in the world. I think that what we'll start seeing is new types of apps that are enabled by that large dynamic computing environment that we'll have access to."

But Porter and the other researchers at CNS are not naive to the security breaches that can occur when users are uploading and downloading sensitive data to and from the cloud, often on mobile devices.

"Any time you're moving away from individual computer devices that you control to a model where there's distributed control there will be new threats," he noted. "When you rely on services that are all interconnected, it becomes really challenging to make sure your data and systems are accessible when you want them, but not when others are

after them."

The pitfalls of distributed computing are not lost on cyber criminals, either. A presentation on "Bitcoin and Cybercrime" by CNS research scientist Kirill Levchenko described how compromised desktop PCs are being used to mine (or computationally generate) Bitcoin. Bitcoin is a decentralized virtual currency that can be easily transferred through a computer or smartphone without an intermediate financial institution. According to Mt. Gox, a major Bitcoin exchange, the price of a Bitcoin in U.S. dollars is more than $100 today, up from just $13 at the start of the year.

A unique feature of Bitcoin is that it can be mined on any computer. Levchenko noted that this makes it particularly attractive to botnets, or collections of Internet-connected programs that can perform tasks. "Bitcoin mining on commodity PCs is not now as cost-effective as it used to be, because the expected revenue per machine is one cent per day, and electricity costs 35 cents per day," said Levchenko. "But with botnets, your electricity is free. If this sounds like printing money, it basically is." Levchenko cited one botnet that was able to earn about $250 per day by mining for Bitcoins.

Levchenko and his colleagues at CNS have also been tracking global Bitcoin use, both legitimate and not. A technique developed by CSE graduate student Sarah Meiklejohn allowed the team to link Bitcoin transactions to major Bitcoin merchants and services, giving the researchers a better understanding of the Bitcoin economy. They discovered that most Bitcoin transactions are to or from exchanges like Mt. Gox. Mining and low-value gambling were the next most common types of transactions. "The success of Bitcoin depends on its adoption for legitimate commerce," Levchenko explained, "but what we're seeing is that it has not yet evolved to that level."

Given the potential for hackers to abuse the cloud, Porter acknowledged the tension that exists between those who want to rely on the cloud for greater computing power and efficiency, but also want their systems and data to be 100 percent secure—especially in light of the recent admission by the U.S. government that it has access to its citizens' digital data.

"That's the problem: If we secure our data too much we lose a lot of the benefits of the cloud, but if we open it up too much we're relying on everyone else to, in a sense, protect us," he noted. "There's a technology angle to this problem but there's also a human angle that as a citizenry we're going to have to get involved and address."

Provided by University of California - San Diego