InformationWeek
THE BUSINESS VALUE OF TECHNOLOGY

Union Pacific Delivers Internet Of Things Reality Check

U.S.'s largest railroad uses sensors and analytics to prevent derailments, but it also shows where the next wave of innovation is needed.

By Chris Murphy, InformationWeek
August 08, 2012
URL: http://www.informationweek.com/global-cio/interviews/union-pacific-delivers-internet-of-thing/240004930

Union Pacific, the nation's largest railroad company, is a choice place to assess the gap between the dream and the reality of what's commonly called the "Internet of things." Like a lot of technology movements, the Internet of things is easy to describe but hard to execute. It means putting sensors on all manner of machines to collect data, linking them over wired and wireless networks, and then applying data analytics to assess when a train's wheel needs replacing, a power plant needs fixing, a soybean field needs watering, or a patient needs reviving.

Union Pacific and other companies are doing some of that kind of analytics today. Companies that succeed in the future will have to do much more of it, so it's critical to know the barriers and constraints--which brings us to Union Pacific headquarters in Omaha, Neb.

UP has been in the same business for 150 years, hauling railcars full of coal, grain, manufactured goods, and people across the plains and mountains of the Western U.S. To run the railroad of the future, it's putting more emphasis on a new skill: predicting the future.

For example, UP would like to be able to predict when a wheel is going to fail weeks before it causes a 1.5-mile-long, 20,000-ton coal train running at 70 miles per hour to derail, risking lives, causing delays, and losing the company money. UP has been using technology to predict and prevent derailments for well over a decade. It has placed infrared sensors every 20 miles on its tracks to take about 20 million temperature readings of train wheels a day to look for overheating, a sign of impending failure. It has trackside microphones to listen for growling bearings in the wheels. For its biggest and heaviest coal trains, it even shoots ultrasound images--like those used to see a fetus in the womb--to look for flaws inside wheels.

BY THE NUMBERS

20 Million

Number of daily pattern matches UP does of sensor readings in order to spot problem wheels

10 Billion

Number of transmissions its firewalls rejected in 2011

$35-$40 Million

Annual revenue UP generates from its own innovations

Data is sent via fiber-optic cables that run alongside UP's tracks back to its Omaha-area data centers. There, complex pattern-matching algorithms flag the outliers, letting experts decide within five minutes of taking a reading whether a driver should pull a train off the track for inspection, or perhaps just slow it from 70 to 35 mph until it can be repaired at the next station. Using all of these technologies, UP has cut bearing-related derailments by 75%.

But the railroad isn't content. Slowing a train to half speed slows the entire network, since the trains behind it on the tracks must slow as well. Stopping a train at a station midroute is a lot better than stopping it out on the track, or having a derailment, but it still bungs up the network.

CIO Lynden Tennison would love to someday get rid of all of those wayside sensors and instead put tiny sensors, called motes, directly on the wheels, bearings, door locks, and more, and collect the data wirelessly. That way, UP could predict weeks out when most repairs are needed and do them when a train is between deliveries and at a major hub. But there's a long list of obstacles to achieving this dream network:

Wireless coverage isn't everywhere. Sensors are too expensive. Unlimited data can't be instantly analyzed, even by today's computers. Algorithms can't replace human judgment. Location data is incomplete. Batteries are too short-lived. "Those are all things that will come," Tennison says.

Until then, he and his team of 2,000 IT pros keep pushing today's technology to make UP a smarter railroad, testing the limits of the Internet of things as they go.

Technology Alongside Judgment

To grow, a railroad needs to push more freight volumes through a relatively finite network of tracks, locomotives, and railcars. IT is essential to maximizing that throughput; it's a constant balance between analytics-driven automation and human judgment.

Railroads are a growth business. After weathering the recession and a disastrous 2009, when revenue fell 21%, UP has climbed back, with revenue up 15% in 2011 to $19.6 billion. The country's seven big railroads have been benefiting from power plants' demand for coal, China's demand for U.S. crops, and recovering auto sales. Even e-commerce growth provides a lift: UP has a section of its dispatch center devoted to one customer, UPS, which is riding the growth of e-commerce home deliveries. Meantime, UP is looking to become more of an end-to-end player by partnering with local trucking and logistics companies.

But UP's core is capacity-constrained. It has one network of tracks, 32,000 miles of it, and adding new track costs about $2.5 million a mile. One new locomotive costs $2.2 million. "We can't spend enough money on capital to fulfill our growth," says Tennison, even though UP's capital spending in 2011 was $3.26 billion. "We have to get there with IT"--by steps that improve decision-making through better information.

Nearly every technology project at the railroad can be tied back to one metric: velocity.

Velocity is the amount of freight moving through the system at a time, and dispatchers talk about the track network like a giant pipeline. UP has about 3,350 trains traveling its tracks and rail yards in any given 24-hour period, and the dispatchers have to manage those trains and their crews to optimize the flow.

Mission control is UP's main dispatch center in Omaha, something of a cross between an air traffic control room and a stock exchange. LED monitors display real-time information on everything from track and train problems to train speeds and locations to fuel consumption to freight counts (the arrangements and contents of shipping containers) to adverse weather patterns specific to its exact track locations to the status of repairs in any one of the company's 120 rail yards and terminals.

To understand the constant blend of technology, automation, and human judgment, consider all of that data coming in from trackside wheel sensors. UP's custom-written software assesses more than 20 million wheel readings a day, generating about 1,500 daily alerts that employees look at. Most problems are considered "strategic"--they aren't a looming risk and just need to be addressed at some point. Fifty to 60 can wait to be checked at the final destination, 25 to 30 at the next terminal, and seven to 10 are serious enough that the driver needs to pull the train over as soon as possible for inspection.

If a train has to slow down, dispatchers must manage the ripple effect throughout the network. An executive VP of operations monitors current and projected velocity constantly, and if it's projected to decline will isolate which trains are dragging it down. And this system doesn't just monitor UP tracks--60% of company shipments spend part of their journey on another railroad's tracks.

Other railroads are doing similar wheel-sensor readings, and ideally this trending information gets shared among the railroads. Since 2006, the major U.S. railroads have had a consortium that works out how to most effectively share that data. Without that data, when a train comes onto the

UP system, analysts can't tell if a wheel suddenly got much hotter or had been steaming along that way for miles.

Another trade-off between automation and judgment comes with the Optimizing Traffic Platform, decision-support software designed using an algorithm from Carnegie Mellon to optimize railroad throughput. It can be configured to maximize throughput, to maximize velocity while giving preference to certain shipments, or to stick as closely as possible to the schedule. UP was the first major railroad to use the OTP technology, in 2010, and it's now used for about 60% of the railroad's operations, with the goal of 100% within a couple of years. Without OTP, the railroad uses a rules-based planner for when trains will meet and have to pass. That planner lets people make the best local decision, but it doesn't factor in the ripple effect. The computing capacity the OTP software requires is one reason UP recently upgraded its long-running VAX 700 data center servers to clusters of Hewlett-Packard and Dell blade servers.

Imperfect Data

The U.S. railroads are proud of their rising efficiency. U.S. rail freight rates are half of what they were in 1981 in inflation-adjusted dollars, the Association of American Railroads is fond of noting. Railroads move a ton of freight 484 miles on a gallon of fuel, up from 235 miles in 1980. Technology plays a role there as well: UP has been testing a system on its Wyoming-to-Chicago coal trains that tells a driver when to throttle up or down to maximize fuel savings.

And yet this complicated network runs on imperfect data. As companies consider the future of analytics, one of the big limitations still is getting the needed data at an affordable price.

For railroads, location data is a big gap. While UP has GPS tracking on several thousand locomotives, GPS isn't reliable enough for safe dispatching and switching of trains--it isn't accurate enough to tell which of two parallel tracks a train is on. So the tracks have a location sensor about every 15 miles that tells a central dispatcher where a train is, leaving a lot of blind spots. And even at those markers, the system knows only that a train passed the spot. In terms of where a train will be in the next five minutes, one traveling 60 mph and accelerating is a lot different from one going 40 mph and slowing down.

Another reason trains don't use GPS is because there isn't universal cellular coverage along the tracks, especially in the vast expanses of the Western U.S., for trains to continually transmit their precise locations. Until the trains get that precise location data and ubiquitous wireless data coverage, railroads can't do a lot of the sophisticated routing, monitoring, and automated controls envisioned for the Internet of things.

"If I could wave my magic wand, I would love to have 4G networks everywhere across my tracks," Tennison says. "That would really change the game for us." The industry is funding research on space-based navigation options like GPS, hoping to find something that meets its requirements.

Sensors are another linchpin to the Internet of things. They're how we'll extract the data from all of these devices. Tennison can quickly spin out a list of places on a railcar he'd like to put sensors: on bearings to measure vibration to anticipate a failure; on a tanker valve to measure pressure and know if there's a leak of toxic gas; on doors to know if a high-value container has been opened. But it would cost UP $500 to $600 to put eight to 10 battery-powered sensors on a railcar. The cost needs to get closer to $200 to be practical. "I think that's coming, but it's coming slower than I wish it was coming," Tennison says.

Those sensors would need to last on a railcar for five to seven years without replacement. That rules out today's battery-powered sensors and likely will require "energy harvesting"--techniques such as capturing the energy from a train braking to use as power.

And UP would need to persuade the industry to standardize on sensors, network standards, and messaging.

Another challenge has to do with how quickly companies can do analytics on the reams of data they're collecting. To understanding that challenge, it's worth a quick detour from railroads to power plants.

The Industrial Internet

General Electric's moniker for the Internet of things is "the industrial Internet." For example, it has a center in Atlanta that remotely monitors about 1,800 GE gas turbines used in electricity power plants worldwide. Sensors on the turbines relay performance data so that GE can anticipate maintenance needs and try to avoid breakdowns.

Software is essential because GE isn't really selling a gas turbine; it's selling the ability to generate power.

GE is increasingly selling service contracts that are less about making repairs and more about guaranteeing performance. It has a pipeline of $147 billion worth of performance-based service contracts that will generate $45 billion in revenue this year, CEO Jeffrey Immelt said in the company's last annual report.

Like at UP, however, what's just as interesting is what GE can't do now but expects do in the coming years.

Until about three years ago, GE had to be selective about the data it collected from those gas turbines. It could collect only about three months' worth of data--about 50 TB--because its database and analysis tools didn't scale beyond that size. See the problem if you were in the energy business? It meant you couldn't trend the current heat wave against the last few years' heat waves. What's more, the software could analyze data from only the turbine itself, not data points from the steam turbine and heat-recovery equipment around it, which also might have clues to a pending breakdown.

Over the past few years, GE has scaled its software so it can collect and use data spanning many years and more equipment. It also acquired analytics vendors such as SmartSignal that let it handle more complicated data, like the effects of those machines connected to the gas turbine.

The next big obstacle GE's software faces is speed--which it will need to predict the future. Today's software can send alerts when it spots a potential problem, but it isn't nearly fast enough to do "what if" analysis of machines. "People call us up and say, 'Can we overdrive our equipment for the next two hours?'" says GE software CTO Rich Carpenter. Today, those answers are given based on decades of experience and engineering knowledge.

What GE wants to offer is the ability to ask, "Has any machine in our entire system ever had X, Y, and Z factors, and what happened four hours later?" GE's systems today would take about 30 days to answer that question--if they could even answer it. GE's working to combine its data management and analytics software with Hadoop-based data processing to deliver an answer in 30 seconds. GE Software has just tested a prototype architecture that delivers that kind of speed, says Erik Udstuen, business leader for GE's software and services business.

"It's being able to predict the future by having a very clear line of sight to the past," says Jim Walsh, global general manager for GE software and services.

Yet GE sees some of the same limits that Tennison sees at UP: sensors are pricey and networks to collect the data can be spotty. It's one thing to cover the 1,800 gas turbines inside power plants in sensors and collect the related data. It's much harder to do that for tens of thousands of wind turbines spread across often remote expanses of the U.S. and China.

This kind of reality check on the Internet of things is essential. People's faith in emerging technology gets to the point that they start to assume that all data is gettable, that all of it is crunchable to turn questions into answers, like those magical computers that spit out the answer in spy movies.

Union Pacific points to what's possible. Yet at the same time, its goal of driving growth through greater use of analytics, sensors, and networked machines shows how much work still lies ahead.