FREIGHT SCHEDULING

Would you route transcontinental traffic via Lake Superior or the Midwest U.S.? Those are the kinds of questions worked out by railroaders in the service design department. Here's what they do

by Steve Ryan

Transcontinental Canadian Pacific trains could run across Canada via Lake Superior (main photo) or through the U.S. via St. Paul, Minn. (above). Main photo, David Young; above, Travis Dewitz

Have you ever watched a train roll by and wondered where all of its freight cars came from or where they're going? Or why the cars are on that particular train and not another, or perhaps on another rail line? If you have, then you've entered the world of service design.

As North American railroads moved from tonnage-based to schedule-based operations in the 1990s, service design departments started to come into their own. From their origins in the traffic departments spread across a railroad's divisions, today's centralized service design departments create comprehensive plans to move traffic across a Class I railroad's broad network of routes.

The idea of designing train service sounds simple enough: Decide how a railcar from station A will move to station B. However, when you start to look at the thousands of cars a day with different origin/destination combinations moving on a typical Class I railroad, this seemingly simple task becomes enormously complex. The role of service design is to create order by giving each car a trip plan to move from origin to destination. Simplicity, efficiency, speed, and reliability are paramount.

A Canadian National local based in Decatur, Ill., serves shippers along a CN branch between Pekin and Mattoon, Ill. In July 2010, the crew picks up grain cars at Emden, Ill. Steve Smedley

IT STARTS WITH THE TRAFFIC FILE

The traffic file is the foundation of the service design department. It contains detailed records of all the traffic a railroad moves — origins, destinations, commodities, shippers, and more — organized into a giant spreadsheet. Each record is a history of car movement. That record might represent an average of fewer than one car a day (such as a shipment that happens once a month) or 50 cars a day or more (as from a coal mine or grain elevator). The traffic file could contain records for upwards of 12,000 shipments a day between 10,000 origin/destination pairs from 5,000 different customers on a typical Class I railroad. To move it efficiently might require 6,000 blocks (groups of cars with a common destination) moving on 500 trains a day.

It's the job of service design to understand the movement of every car on the network and develop plans for its handling. In the case of three boxcars loaded with corrugated paper moving from Calgary, Alta., to Harrisburg, Pa., via Toronto (see below), service designers asked all kinds of questions before developing a trip plan: What if the daily volume of traffic going from Calgary to Toronto was only those three cars instead of the typical 40? Would you still run a direct Calgary-Toronto train,

HOW MANIFEST FREIGHT MOVES

To see service design in action, let's follow a group of boxcars loaded with corrugated paper in Calgary, Alta., and destined for a warehouse in Atlanta. Two railroads will handle this move, with an interchange at Harrisburg, Pa. The paper company loads three boxcars a day, seven days a week. That's not enough volume for a unit train, so the boxcars will move in manifest freight service in tandem with other cars between terminals. This type of service lengthens transit times for the customer, but maximizes the railroad's economies of scale, keeping costs and rates down.

Cars with common destinations move grouped together in blocks, and every block is assigned to a train. Our boxcar loads will move in blocks with other cars for their entire journey, beginning with the local freight that picks up the boxcars from the customer — one of 10 shippers that local freight serves. At the end of the day, the local brings the cars it has lifted to a hump yard at Calgary. The cars roll down the hump and enter a bowl track, where they are grouped with other cars all headed for the same place, thus creating a new block.

The new block for these cars is Toronto. The block averages about 40 cars a day, and is assigned to a mainline train that leaves for Toronto at 3 p.m. daily. However, a 40-car train is too small to maximize efficiency on a road freight traveling 2,100 miles, so the train also gets a Winnipeg, Man., automotive block consisting of empty auto racks that averages 20 cars a day. And on the way out of Calgary, the train stops at an intermodal terminal where it picks up double-stack cars also bound for Toronto. To make them easier to set out later on, the crew attaches the stack cars on the train's tail end. Now the train has 120 cars in four distinct blocks: auto racks for Winnipeg, mixed freight for the Toronto hump yard, and two intermodal blocks — one loaded with containers and one of empty flatcars being repositioned.

At Winnipeg, 830 miles east of Calgary, the train sets off its auto racks and adds a Winnipeg-Toronto manifest block to the head end, creating a three-block train (Toronto mixed freight plus the two intermodal blocks). The train continues on to Toronto, stopping only to change crews, refuel, or meet other trains at sidings. At the Toronto intermodal yard, the train crew spots the block of loaded container cars directly on a working track, where cranes can get at them. The empty cars go on another track to be loaded later. The rest of the train then heads to the Toronto hump yard. Like other terminals, it builds dozens of blocks a day. The boxcars of paper roll into a classification track joining other cars for Harrisburg and the connecting railroad.

In their 2,600-mile journey from Calgary to Harrisburg, the cars will move in three blocks: a local, a manifest, and an interchange block. Why three blocks and not five? Why three trains and not two? Those are questions the service design department analyzed before determining that the three-block move was the most efficient for the rail network and the customer. — *Steve Ryan*

This CSX unit train is operating on Norfolk Southern at Copper Creek, Va., though a CSX bridge soars above. Frank Lassiter

or would you send the cars to a yard in Moose Jaw, Sask., mixed with other cars bound for Brandon, Man., and Estevan, Sask.? Would Moose Jaw then have enough traffic to build a Toronto block for a train? If not, could the boxcars move in a block to Chicago and reach Harrisburg that way? For railroads, economies of scale are critical. Without them, costs soar. The fewer blocks a yard has to make, or the fewer blocks on a train, the easier it is for the operating department to implement the plan, with fewer chances for errors.

There are limits, though. If you reduce the number of blocks per train, more yards will likely be sorting more traffic. The more yards that touch a car, the greater the cost and chance for delay. It takes longer for cars to reach Toronto if they have to stop in Moose Jaw for reclassification versus running through. Every yard has a set capacity of cars it can process and blocks it can build in a day. When trip plans send more cars to a yard than it can handle, service designers have to look for opportunities to redirect traffic to less congested terminals. A railroad's book of business — its mix of traffic — will dictate whether those boxcars should make their journey via Winnipeg, Man., in three blocks or via Chicago in five blocks.

UNIT TRAIN CHALLENGES

Unit trains might seem easier to design, since the train is essentially one block moving from point A to B. But they pose a different set of challenges. Among the first questions service designers ask is whether a commodity move fits the unit-train definition. Will the traffic be able to travel in a continuous cycle, seven days a week? Can the locomotives remain attached during loading and unloading to eliminate costly positioning moves? Can the train be loaded and emptied in 24 hours or fewer? Can the unit train's length be maximized to match the capability of the network? Is the traffic time-sensitive, and what would the car-cycle time difference be between unit and manifest service? Can the originating point build a unit train, and can the destination handle one? If not, are there places to store and switch the cars at each end? Are there enough crews and motive power

BNSF'S 1990S MANIFEST NETWORK

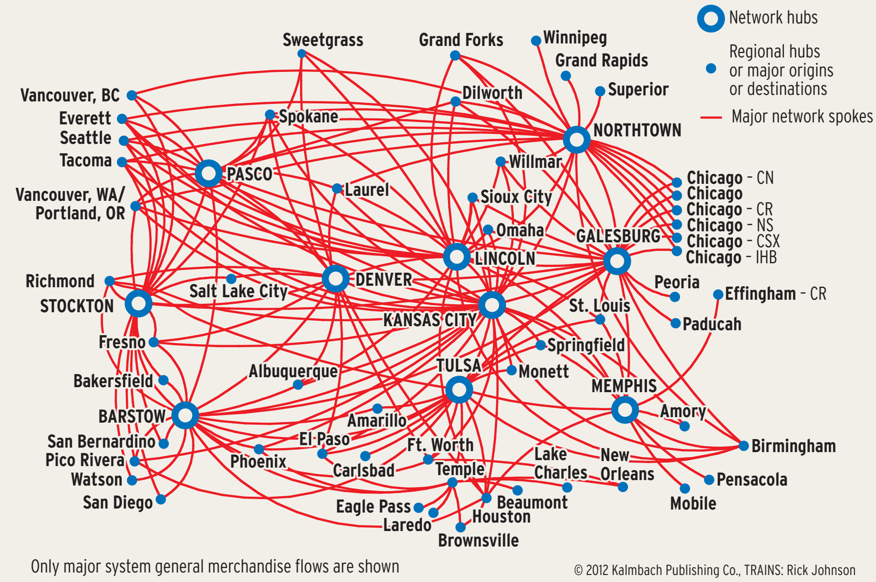

BNSF Railway developed this diagram of its manifest freight network in the 1990s, although the principles remain the same today. Freight cars, like airline passengers, move in a hub-and-spoke pattern, making connections at regional and network hubs to minimize handling and optimize volume. In 2001, BNSF's carload network had 13,400 origin/destination pairs, although 90 percent of the shipments moved between just 5,000 origin/destination pairs.

HOW DOES A CAR KNOW WHERE TO GO?

Before a car can be assigned a trip plan from the service design department, it has to have a destination. For loaded cars and privately controlled empty cars, the process is fairly straightforward. Not so for railroad-controlled empties.

Loaded cars inherently know where to go. Shippers send the railroad billing information detailing a car's destination, receiver, routing (interchange points and other railroads), contents, and more. Historically, the shipper submitted this information on a bill of lading for the railroad's local freight agent. The agent would then create a waybill that would physically move with the car, carried by the conductor. Yards used the waybill information to determine how to switch the car. Lose the waybill and you had a "no-bill" — and no idea where the car went.

With computers, local agents were replaced by centralized service centers where customers faxed their bills of lading. The shipping information went into a computer system, from which yard employees could print out waybills. In the 1980s and '90s, railroads rolled out electronic data interchange and customers began transmitting the bill of lading data electronically. The conductor's train list (a printout of all cars in a train, with customer and shipping information) replaced the stack of waybills.

Empty cars fall into three basic categories, each with its own procedures. Privately owned or leased cars (called "X" cars because their initials end in X) generally go back to their last loading point via the reverse route, unless the railroad has instructions to the contrary.

The "Car Service Rules" administered by the Association of American Railroads determine the disposition of foreign cars, which are railroad-controlled cars that do not belong to the railroad they are on. (A Norfolk Southern car on Union Pacific is considered foreign to UP.) Depending on the car type, whether it's assigned or unassigned, and whether the car belongs to a direct connection or an indirect connection, the rules will dictate how it should be handled when empty. The car may go back the reverse route, to the home road at any junction, or follow a specified route if the home road is not a direct connection. Two supplemental aids in this process before automation were the stenciling on the sides of cars ("When empty return to …") and the waybill, which often contained empty disposition instructions. Now when cars are reported empty, the computer system automatically determines a way back home and assigns it a new destination.

When cars are bound for interchange, railroads often send the receiving carrier a train list in advance. If the delivering railroad errs and disposes of an empty car improperly, the receiving carrier can refuse it and send it back as an "error move." In pre-mechanized days, a phone call from the receiving railroad would alert the delivering carrier of the car's return and the erring road would have a proper waybill waiting. Now, with less human attention given to individual car movements, it is not unusual for a car to play "railroad ping-pong": Railroad A gives the car to Railroad B in error. Railroad B's computer recognizes it is not obligated to handle the car under the rules, so it returns the car back to Railroad A. When Railroad A receives the car back, its computer thinks it is not obligated to handle the car from Railroad B under the rules (not recognizing that it started the chain), so it refuses the car back to B. And so it goes, until someone happens to notice and intervenes.

System cars are the railroad-controlled cars that are "home" (a BNSF Railway car on BNSF), and their next destination is determined by the railroad's car management department. This group manages the car fleet and ensures that customers receive cars for loading by the required day. Before computers, the process was described by one Chicago & North Western car management veteran this way: "Order, counter-order, disorder, confusion, chaos. It amazed me anything got anyplace. There was no instant communication. We never had any idea what was happening." The central office would issue directives (such as "50-foot boxcars at Chicago go to Green Bay [Wis.]") to car distributors located on each division, but cars would often get siphoned off while en route (in places like Milwaukee) and many never made it. The division car distributors would take car orders from the local freight agents and provide disposition instructions for

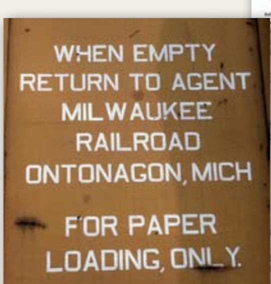

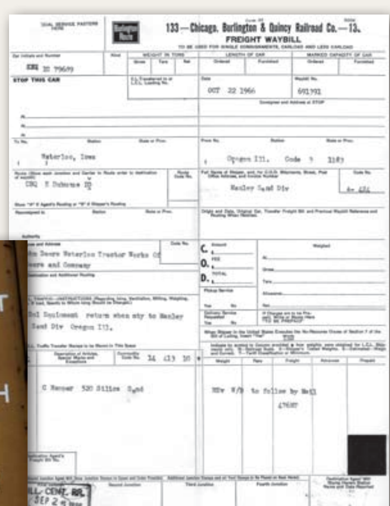
133—Chicago, Burlington & Quincy Railroad Co.—13.
FREIGHT WAYBILL
STOP THIS CAR
OCT 22 1966
Waterloo, Iowa
C.
O.
D.

Historic car management tools include orders stenciled on freight cars (left) and waybills (right). Two photos, Brian Buchanan

CSX trains with UP and CP locomotives prepare to depart CSX's hump yard at Willard, Ohio, in 2012. Brandon Townley

HOW MANY CARS A DAY DO THE BIGGEST RAILROADS HANDLE?

Daily U.S. car counts for the week ending June 19, 2012

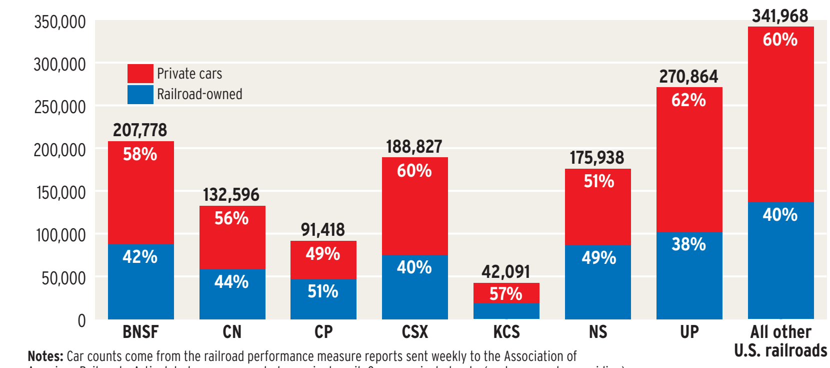

Notes: Car counts come from the railroad performance measure reports sent weekly to the Association of American Railroads. Articulated cars are counted as a single unit. Cars on private tracks (such as a customer siding) are counted on the last railroad on which they were located. Counts do not include maintenance-of-way cars.

empty cars. Foreign cars would be disposed of per the rules and system cars not needed elsewhere on the division would be flowed per the directives from the central staff.

Now, each major railroad has its own computerized car distribution system and centralized staff. Similar car types are usually grouped in the system as pools. (For example, all 52-foot, 100-ton, plain-interior open gondolas with sides ranging from 4½ to 5 feet are in one pool.) The system matches up available cars in a pool with customer orders entered into the system specifying car type needed, quantity, loading destination, and date required. These orders can be set up to draw cars from only a single station, a region, or from the whole railroad. A car reported empty (whether released from an industry or received in interchange) will trigger a search by the system for an unfilled order meeting the car's criteria. Orders can be ranked in hierarchy or prioritized by loading date. Cars not hitting a customer order usually drop down to a "flow order" and are directed to a redistribution point or storage location.

Car management employees monitor car orders to ensure they are being filled, as well as the railroad's car flows, making manual changes in the system as necessary. For instance, if a block of cars moving to a loading point misses a train connection and the trip plan is set back a day, the cars may arrive too late to meet the need. In that case, a different set of cars may have to be found that can fill the order on time, and the delayed cars will get a new destination.

— BRIAN BUCHANAN. The author has worked for four railroads during the past 31 years in operations, car management, and marketing.

to handle this business as a unit train? Could an existing scheduled service accommodate the traffic instead of a dedicated unit train? What does the market require? These are just some of the questions the service design department considers.

A Class I railroad may have several routes a unit train could take. Service designers will evaluate the advantages and disadvantages of each. Sometimes the choice is simple: Use the most direct mainline path. However, factors like corridor capacity — the number of trains per day a main line is rated for — may dictate a different route. Consider the possibilities for a 100-car unit grain train moving from an elevator in North Dakota to the port of Montreal. Two routes are possible. The first is the natural flow, which would be southerly via St. Paul, Minn., and Toronto. The second one is slightly longer, and sends the train in a northerly arc via Portal, N.D., Thunder Bay, Ont., and Toronto. In addition to line capacity, planners also factor in each line's ruling grades. Would two AC4400CWs on the head end pulling a 13,000-ton unit train surmount the line's ruling grades, or should the railroad consider a distributed power model? Service designers calculate the costs for each route, and identify the pros and cons of each.

Designers may decide that the longer northbound route, although more expensive to operate and needing an extra crew, is the way to go. If capacity constraints on the shorter route drove that decision, the grain train's northerly routing could be just a temporary solution. The new traffic may prompt the railroad to invest in new sidings along the southerly route, install centralized traffic control, or even redirect other trains.

THE INTEGRATED DESIGN

For both unit and manifest trains, service designers must consider mainline and yard constraints when developing schedules. It does no good if trains are bunched

Motive power is a factor in service design. Mid-train distributed power units assist a UP coal train at Rollinsville, Colo. Travis Dewitz

together and all arrive at a yard within a short time frame. Tracks will quickly fill up, and the wait to get locomotives serviced and back out could become needlessly long, wasting resources. Spacing train arrivals and departures throughout the day lets yardmasters better carry out a terminal's work, although that spacing isn't always possible.

For example, Southern Ontario has an abundance of auto plants and intermodal facilities. When the workday is over, all those new automobiles and all those containers that were trucked in need to move out to make room for the next day's production. That means that nighttime is prime time, and the mainline trains have to work quickly to pick up the day's traffic and make room for the next day's shipments.

In general, it's easier to schedule trains departing from an originating yard. Local terminal superintendents advise what windows are best for building a train, based in

A CP local returns to Wolverton Yard on July 11, 2012, with auto racks loaded at a Toyota plant in Cambridge, Ont. Greg McDonnell

SERVICE DESIGN IN ACTION: PLANNING A NEW YARD

Increasingly, service design is playing a role in planning railroading's new infrastructure. Canadian Pacific's new yard in Wolverton, Ont. [see "News," February 2012 Trains], between Toronto and London, Ont., provides a good example. Opened in 2011, the yard stages shipments to and from three nearby auto plants. CP's service design department, where I work, was part of the planning team from Day 1.

First, we looked at the area's existing infrastructure, its volumes, how many blocks it produced, how many trains were involved, and where the work was currently being performed. Then we looked at the projected traffic volumes provided by Commercial (sales and marketing), and determined what the new yard would need to support those volumes. Discussions followed with the engineering and operating departments on what the yard's track configuration should be. We explored new yard designs. Something as simple as adding a crossover would prompt more discussions, as we considered the advantages and drawbacks of each possibility. Should the yard be used to switch traffic, or should it function mainly as a block-swap facility? How many tracks should the yard have, and what work could be pulled out of other yards? How many trains should stop to pick up or set out traffic?

The result was a new yard built to support a service plan designed to move traffic in the most economical way possible, while meeting the needs of the shippers. The yard has one 11,000-foot siding and three 10,000-foot yard tracks. Twenty-one trains call at Wolverton Yard (not all of them daily), including 15 mainline freights and two locals based there. From Chicago to Toronto, the service design department altered the trip plans of 34 trains. — *Steve Ryan*

part on how many blocks a train has. Building a two-block train is quicker than a six-block train. Where it gets tricky are the intermediate yards. They are the proverbial meat in the sandwich, and moving trains around to space them out by schedule is difficult. Staging trains en route to reach a yard at a desired time is not optimal either, as this increases transit time and wastes assets.

Service designers also collaborate with terminal managers and other departments to work out the finer details of a train's schedule, such as where to perform inspections, fueling, and crew changes. All those elements need time built into the schedule, and influence what the most efficient way to do the work might be. Commuter train curfews, meets en route, and other potential delays also enter into the calculations.

As Class I railroads have grown in size and complexity, the value of having a scheduled service has increased. The engineering department can plan work blocks ahead of time by looking for natural windows between scheduled trains, or find times that will impact low-priority trains but keep the time-sensitive freight moving. Locomotive management understands the power demand by corridor and yard and can plan locomotive servicing accordingly. Longer term, management can evaluate the need for new locomotives. Car fleet sizing depends in large part on how quickly freight cars can be cycled, which in part is a function of a scheduled transit.

Another factor service designers consider is balance, which is created by running daily scheduled services, with sister trains in the opposite direction. This reduces nonrevenue moves of crews or equipment, and keeps cars cycling where they need to go. By running the same plan day in and day out, railroads can maximize the efficiency and quality of their product.

CHANGING TRAINS AND DESIGNS

A train's design is constantly changing. In one way or another, most big railroads tweak their network designs on a weekly basis. What can prompt a change? Shippers may grow or shrink their volumes. Contracts may be won or lost. New opportunities can start up or take off. Mergers and acquisitions will prompt more dramatic service changes, as would a large infusion of new business, such as the unit trains of crude oil that began moving out of North Dakota in 2010. Whenever changes occur to the physical plant in a yard or on the main line, designers alter the service plan to take advantage of the new infrastructure. Inevitably, one change will create a domino effect, requiring other train schedules to adjust as well.

Service designers are constantly reviewing the railroad's traffic file, looking for trends and opportunities to make the schedules better. Much of this work is done using sophisticated computer models that analyze and dissect traffic based on algorithms [see "Railroads Meet the Wizard," November 2010 Trains]. However, designers also benefit from manually reviewing train consists, counting cars, and looking for anything unusual or missing.

How long does it take to make a design change? That depends on how complex and extensive the changes are. Some schedule adjustments can take a matter of hours to implement, such as changing the classification codes to direct traffic away from a track washout that will take time to repair. Conversely, a new customer shipping unit-train loads might require several months to work out the optimum train size, power configuration, and routing.

The recession that began in December

New business creates new design needs. The Black Thunder West loadout opened in 2008 on the joint BNSF-UP line in Wyoming's Powder River Basin coal fields. Travis Dewitz

Westbound BNSF and CN freights cross at Berwyn, Ill., on March 13, 2008. Ben Muciek

2007 created a different set of service-planning challenges. The traditional seasonal traffic peaks never materialized, and as the financial crisis worsened, volumes started to nosedive. Service designers scrambled to understand what it meant. As traffic levels fell, costs had to decrease as well, prompting train design and traffic flow changes. From a service design perspective, patterns that move traffic consistently throughout the year, rather than in peak and lull cycles, produce better service. Planners used to make significant scheduling changes during the spring and fall peaks, when thousands of extra cars would flood the network. Now, service designers are wondering when, if ever, those peaks will return. (So far they haven't.)

WHERE SERVICE DESIGN FITS IN

The size of a service design department varies. At a short line, the railroad president (who might also be the locomotive engineer) may likely perform this work. At a regional railroad, often the chief of transportation or superintendent will handle the scheduling responsibilities. They understand the flow of the whole network and its abilities, and they often can quickly pull together sales, marketing, and operations people to work out a service plan.

At a Class I railroad, the service design department can consist of dozens of employees, each managing one piece of the puzzle. One person might handle train profiles, another might maintain the background software systems, and still others will act as the primary contacts with operations or sales and marketing. To maintain a big-picture focus, the department usually has a few individuals, similar to architects, whose responsibility is the whole picture or network. They can combine input from other department members to create a unified and cohesive service design. They also help establish policies to meet the goals set by senior management.

Most service design employees have a background in operations, transportation, or sales and marketing. Service designers need some understanding of operations to know why a requested schedule change might not work. Likewise, a marketing background is important to understand a customer's transit requirements. Backgrounds in other fields, such as finance or labor relations, can also be helpful. Service designers must strategically analyze a problem, understand the role of costing, think analytically, and think from a network perspective. Service is the only product railroads have to sell. It's the responsibility of the service design department to provide the best service plan for the railroad and its customers. I

STEVE RYAN has worked at Canadian Pacific for 18 years, eight of them in service design. A Montreal native, he lives in Calgary with his wife and two children.