Modeling the Staffing Boom-and-Bust Cycle to Identify Meaningful Measures and Improvements

Dart Lindsley and Dawn Chaney
Senior Business Process Architects
Cisco Systems, Inc

Most of us have been through something like this: our company starts to ramp up hiring. At first the recruitment staff can handle it, but soon starts to fall behind headcount targets. As a result, the company hires more recruitment staff to meet the demand. This takes a little while. In the meantime more requisitions open up, overwhelming the staff's ability to hire. The growth in the number of requisition appears to be accelerating faster and staffing activity builds to a frantic pace. Then, just when things are looking their worst, headcount target is suddenly close; in fact, the company is approaching headcount very rapidly and not slowing down. Finance reacts, slams on the brakes, and cancels most of the requisitions. This is the staffing boom-and-bust cycle.

Is "boom-and-bust" good business? Not during the boom: the company increased headcount because there was an opportunity to be had, money was on the table, and the company was in a race with competitors to pick it up. The company suffers the opportunity cost of being unable to hire enough employees to take that prize. And it's not good business during the bust either. When requisitions are canceled, most of the recruitment staff is idled and may have to be let go – and work on canceled requisitions is wasted.

Why does this rise and fall in staffing workload occur? Is it unavoidable? Is it our fault? We assert that the way companies maintain headcount is a dynamic system-a system in which causal feedback loops can result in severe fluctuations. We suggest that modeling the cause-and-effect relationships within this system, using systems dynamics modeling techniques, can help us to answer these and other questions. We demonstrate how to create and analyze such models to identify which measurements can be taken and which levers can be pulled to reduce the severity of the staffing boom/bust cycle. We conclude by discussing briefly how similar models have been used at Cisco Systems to target the specific strategic process improvements of measuring more frequently and reducing time to fill.

Example 1: Basic Model in Steady State

Figure 1 is a simple model that employs the basic building blocks of system dynamics modeling to describe how a company grows or shrinks its employee population when the headcount target is held constant. The rectangle labeled Employees, called a stock, represents all the employees in the company. It behaves like a bathtub that can be filled by the flow of employees coming in, labeled Hiring, and can be emptied by the flow of employees exiting, labeled Leaving. The little clouds indicate that the employees are hired from and leave to a place outside the system represented by the model. Each flow is like a valve and contains a formula that determines how many employees flow through it.

The circle labeled Leaving Rate contains a number representing the rate at which employees exit the company. The thin arrows carry information from the Employees stock and the Leaving Rate variable to the following formula in the Leaving flow that determines the number of employees who leave the company:

Leaving = Employees x Leaving Rate

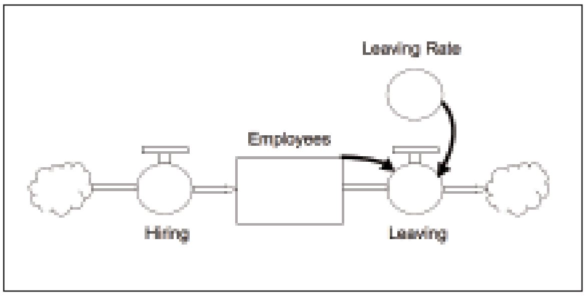

Figure 1: Partial Model: Employees Enter and Leave Company.

For the purposes of demonstration, the size of the company is set at ten thousand and the Leaving Rate is set at 10 percent per year. In Figure 2, the model includes a new stock representing the number of requisitions. Similar to the Employees stock, there is a flow into Requisitions labeled Opening and a flow out labeled Filling. For simplicity, assume that for each requisition filled, one employee is hired. Thus an arrow carries information from Filling to Hiring and the formula for Hiring is:

Hiring = Filling

The formula for Filling is very much like the formula for Leaving. It is a function of the variable Filling Rate:

Filling = Filling Rate x Requisitions

The Filling Rate in the model is set at 30 percent per month. This introduces a delay in filling requisitions that mimics the natural variation in time-to-fill. Now we have everything except a formula for Opening, which is where things get interesting.

The goal of this system is to maintain a stable head-count, a steady state in which Hiring equals Leaving and Employees stays level. To do that, we need to open a number of requisitions sufficient to keep up with turnover. It is tempting to suggest a policy of opening one new requisition as each employee leaves, but such a policy results in the number of employees falling several months' worth of turnover below target. Because of the delay in Filling, it is necessary to open a number of requisitions significantly greater than that of monthly turnover. Figure 3 illustrates why this is: If only a fraction of the total requisitions are filled each month, it is necessary to open enough requisitions such that the fraction equals the number of employees leaving.

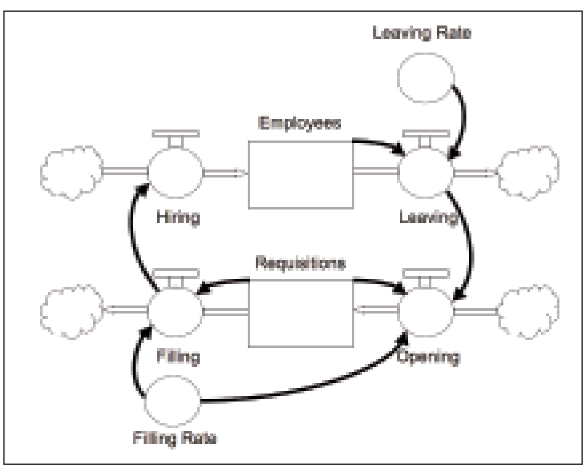

Figure 2: Partial Model: New Stock for Requisitions

Understanding this, it is possible to put the model in a steady state with the following formula for Opening:

Opening = ((Leaving/Filling Rate) + Leaving) - Requisitions

Leaving is added again because if we only open Leaving/Filling Rate we still fall behind the headcount target by one month's worth of turnover. The current number of requisitions is subtracted because we don't want to add this number of requisitions every month; we only want to make up the difference between what is currently in the Requisitions stock and the necessary number. In other words, we are just topping off the Requisitions stock. Figure 4 shows the system in steady state. Both the Employees and Requisitions stocks are flat.

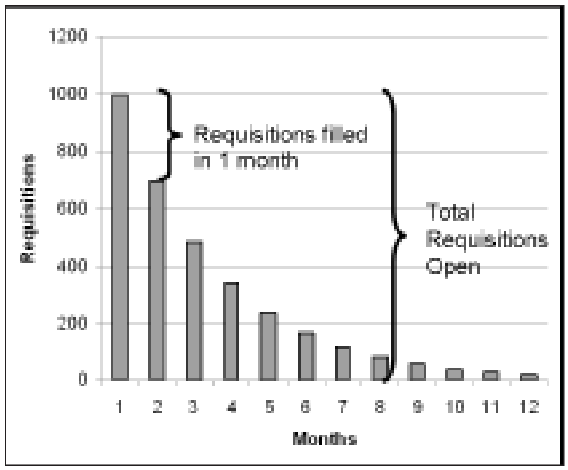

Figure 3: Number of Open Requisitions Each Month Over Time

So far, the model has shown that the longer it takes to fill requisitions, the more requisitions must be kept open at all times to keep up with turnover. The model in Example 1 relies on a number of assumptions to keep it simple. One is that the company's headcount target is constant. Example 2 introduces an increase in the company's overall headcount target.

Example 2: Increase Headcount Target

A headcount target increase is fundamentally

Figure 4: Headcount and Requisitions While Strictly Managing Turnover

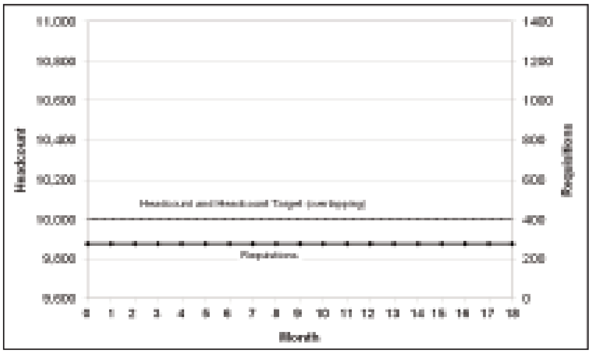

different from keeping up with turnover because it is usually discrete event whereas turnover is continuous. This difference is critical to how the model responds.

The model shown in Figure 2 is modified in Figure 5 to incorporate a variable Headcount Target that provides information to the Opening flow. The formula for Opening is updated to incorporate information from Headcount Target and Employees. For each headcount we are short, open a requisition:

Opening = Headcount Target - Employees - Requisitions

First subtract Employees from Headcount Target to determine the headcount shortfall. Then subtract Requisitions so that instead of opening as many requisitions as the shortfall every month, Requisitions is topped off at a level equal to the current shortfall. Combining this with the formula to keep up with turnover yields an Opening formula of:

Opening = (((Leaving/Filling Rate) + Leaving) + (Headcount Target - Employees)) - Requisitions

These two updates are represented in Figure 5. How this model reacts to the change in headcount is depicted in Figure 6. In this model, a Headcount Target increase causes a spike in Requisitions that tops out at about twice its steady-state value; then over about 20 months the requisition count returns to the values required to keep up with turnover.

This model demonstrates a problem: it would take the company six months to reach 80 percent of the new headcount target, and one year later the company would still be about 3 percent short of target. We can assume that any model that takes more than a year to reach target is unacceptable. Thus Example 3 modifies this model again to reach target more quickly.

Example 3: Accelerate Hiring to Meet an Increase in Headcount Quickly

Let us build the delay in filling requisitions into the Opening formula to meet the Headcount Target more quickly. We can modify the Opening formula so that the headcount shortfall is divided by Fill Rate:

(((Leaving/Fill Rate) + Leaving) + ((Headcount Target - Employees)/Fill Rate) - Requisitions)

The result, shown in Figure 7, is that the company meets the new headcount target within 2.5 months. This meets headcount targets more quickly, but introduces several new problems. For one, the staffing workload spikes to over 430 percent of the number of requisitions necessary to keep up with turnover, then falls to 35 percent over 10 months, before returning to the steady state. Even worse, the company exceeds target headcount by over 50 percent of the increase in Headcount Target (in other words, if Headcount Target increased by 500, then in this model the total headcount would surpass target by 250 employees).

Further, the number of employees remains over target for 9 months before attrition reduces the excess. This is because, to make headcount within a reasonable time, we open many more requisitions than are required and continue filling them after the target is reached. Keeping headcount at or below target is critical to earnings, so we can assume that any model that surpasses headcount by a significant amount and for any significant duration is unacceptable.

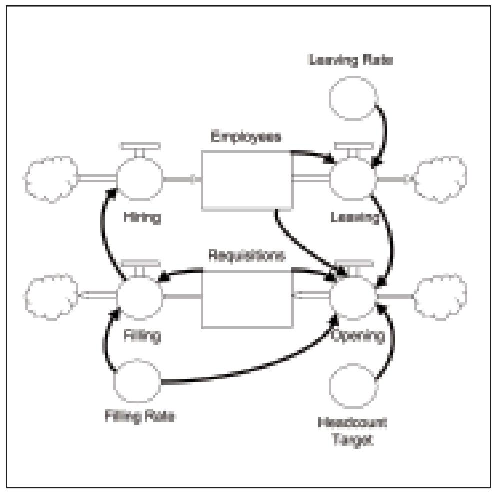

Figure 5: Basic Model with Addition of Headcount Target Variable

We have so far introduced two constraints: the first is that headcount must be met within a reasonable amount of time, and the second is that headcount can't be exceeded. In these models, holding fill rate steady, the outcomes of meeting headcount quickly and not surpassing headcount are mutually exclusive. We can do one, but not both.

In fact there are more cancellations than the total rise in headcount target; that is, if the headcount target went up by 500, this model would result in canceling 520 requisitions. That is the boom-and-bust staffing cycle.

Figure 6: Headcount and Requisitions Given Headcount Increase

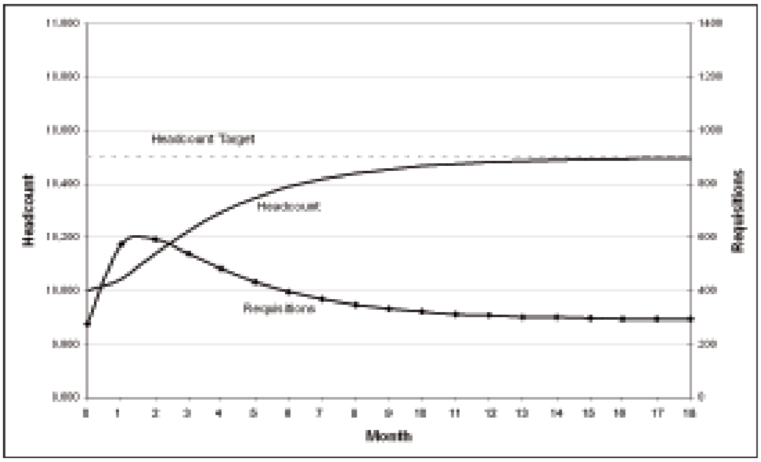

One way around this situation is to provide the option to cancel open requisitions, so that whenever we exceed headcount we cancel and stop filling all requisitions. This requires a few additions to the previous model, as shown in Figure 8.

The formula in the flow for Canceling requisitions is a logical statement that basically means: "If we are over headcount then cancel all requisitions, but if we are under headcount target then do nothing":

IF (Employees > Headcount Target) THEN (Requisitions) ELSE (0)

And we also change the Filling formula so that it means: "if we are over headcount target then stop hiring, but if we are under headcount target then keep hiring":

IF (Employees > Headcount Target) THEN (0) ELSE (Requisitions x Filling Rate)

The result, as shown in Figure 9, functions within the two constraints identified: the company meets the new Headcount Target within 3 months and exceeds headcount by 1 percent of the increase in Headcount Target (and then only for a month).

There is a price to pay for this result: The number of requisitions rises and falls again by 435 percent, but this time in only 4 months instead of the 10 months in the previous model. And even more requisitions must be canceled.

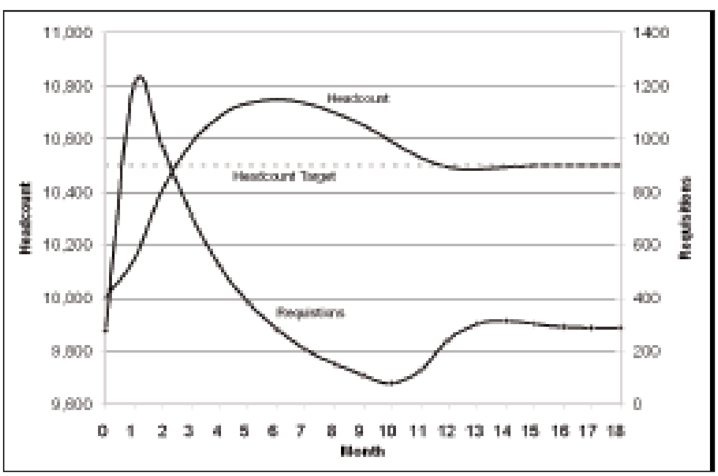

Figure 7: Headcount and Requisitions Accelerated to Meet Headcount Target

Avoiding the Tradeoffs

At Cisco Systems, developing idealized models like the ones described here has provided an understanding of the dimensions of certain business problems in a way that no other tool could have. System modeling has helped us to understand the tradeoffs between meeting headcount targets in time, surpassing headcount targets, and cancellations.

The costs of not meeting headcount targets in time and of surpassing headcount targets are well understood, but the cost of requisition cancellations is often not fully appreciated. First, time and money are wasted on requisitions that will never be filled-in a manufacturing environment, this kind of waste is called scrap and is usually one of the first targets of a process improvement team. Second, the plans of managers who expect additional headcount are disrupted. Third, in this first-come-first-served model, the ability to target which organizations and which skills are staffed is greatly reduced. Fourth, measures of time-to-fill are rendered meaningless because in this model half the requisitions created to meet a new headcount target are filled within a few months, but the other half are canceled before they are filled and don't show up in the time-to-fill numbers at all.

Ultimately we determined that the cost of all three tradeoff options is unacceptable, and simply opting for the least bad is not good enough. We identified two ways to reduce all three. One lever we can pull to avoid these

Figure 8: Basic Model with Addition of Canceling Flow

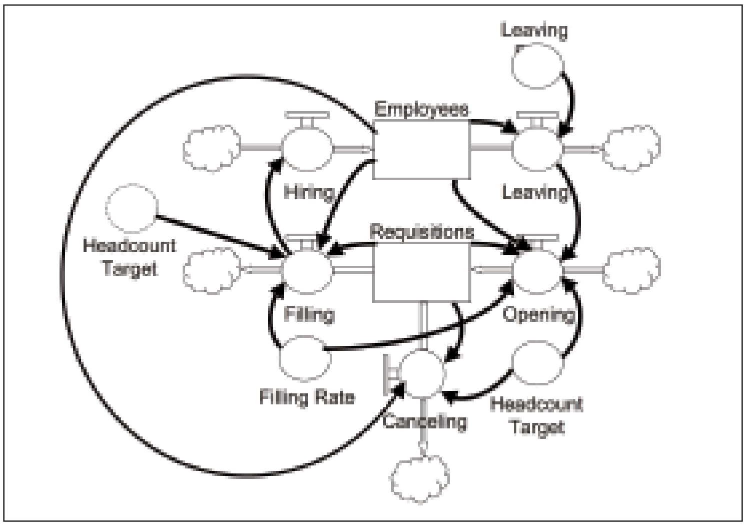

tradeoffs is to measure the system more often. In the models shown, some waste is caused because the time step is 30 days. For example, the system can be one short of headcount at the end of one month and above headcount on the first day of the second month, but it does not respond to the overage by canceling requisitions until the end of the second month. A lot of hiring can happen in that month. So, in general, measuring more often can result in some improvement.

Controlling variation in time-to-fill is truly effective in controlling the severity of the boom-and-bust cycle. It is the delay in the signal from headcount demand to fill that is the principal cause of the inefficiencies in the models. Figure 10 shows how changes in filling rate have a dramatic effect on system performance. These graphs were created by repeatedly running the model described in Figure 8 with different values for the filling rate.

Raising the filling rate from 0.2 to 0.4 decreases the spike in requisitions by about half and decreases the number of cancellations by 75 percent. Improving the filling rate from 0.4 to 0.6 decreases the spike in requisitions by about one-third and decreases the number of cancellations by about two-thirds. Two things are clear from this exercise: there are significant gains to be made by decreasing time-to-fill, but as time-to-fill improves it gets harder to realize gains.

Conclusion

By illuminating more of the true costs of slow hiring, these and similar models have been a powerful tool for advocating on behalf of projects to reduce time-to-fill at Cisco Systems. These analyses have been well received by both HR and Finance executives (who share the issues caused by the boom-and-bust cycle) and have raised the level of dialogue between the two organizations.

Models like these must be used with caution. They can be persuasive and still be incorrect. It is vital to model, then compare to real-world data, then model again.

So far, Cisco's experience and historical data show a correlation with the models in this article. We have lived through spikes in requisitions and cancellations, and we can see similar spikes going back as far as the data go, but our data are incomplete and we are just now beginning to collect it in a way that will help confirm or refute these models.

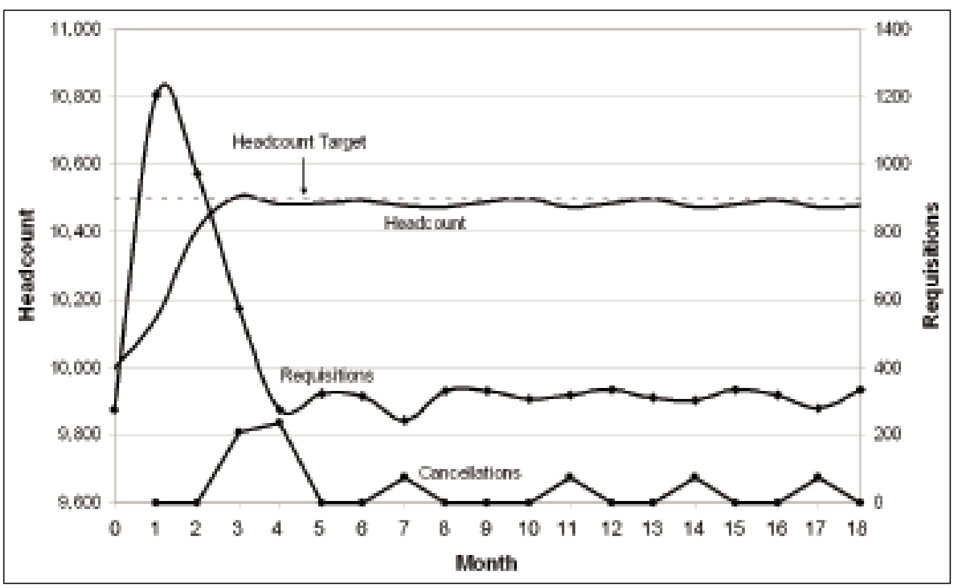

Figure 9: Requisitions Canceled When Headcount Target Reached

Figure 10: Requisition Spike and Cancellations as Percentage of Headcount Increase

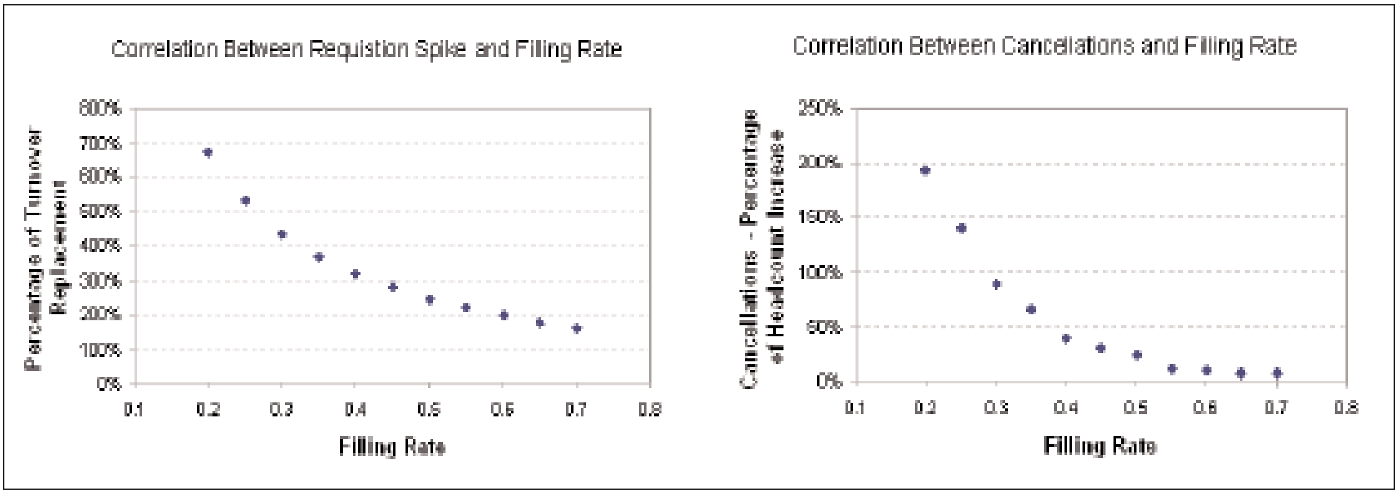

Using system dynamics modeling has, however, greatly improved our understanding and helped us to focus our efforts.

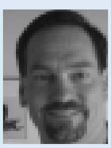

Dart Lindsley
Senior Business Process Architect, HR IT Department, Cisco Systems, Inc
Email: dlindsle@cisco.com

Dart's previous roles at Cisco include Senior Technical Recruiter. Prior to Cisco, Dart was Lead Strategic Recruiter at semiconductor capital equipment maker KLA-Tencor and Staffing Researcher at HR consulting firm R.W. Stearns.

Dawn Chaney
Senior Business Process Architect, HR IT Department., Cisco Systems, Inc
Email: dchaney@cisco.com

Dawn has served in a variety of management and senior analyst roles within Cisco supporting process design and improvement efforts to increase workforce productivity.

Management Tip

"Managing Your New Recruiters"

Provided by Karen Kesner, SVP, Recruiting at Washington Mutual
karen.kesner@wamu.net

CORE (Cool Onboarding Recruiting Excellence) is a training and development program that was recently developed and rolled out for all newly hired Recruiters at WaMu. It was developed in response to a lack of a positive, consistent onboarding and training program in the organization.

The CORE program consists of two parts: "Cool Onboarding" and "Recruiting Excellence."

"Cool Onboarding" ensures that all WaMu Recruiters and Recruiting Assistants (RA) experience a positive welcome to WaMu and are onboarded effectively and consistently, enterprise-wide. The program consists of a number of tools that provides training on the basic, necessary knowledge to manage the hiring process at WaMu. Tools include training guides for both the new employee and current team members responsible for training, a checklist to ensure that the manager doesn't miss any of the crucial details for onboarding the new team member, a job shadow program, and a self-directed e-Learning class that takes them through the WaMu recruiting cycle.

"Recruiting Excellence" is dedicated to the continued training and development of all of the WaMu Recruiting Team. Curriculum is currently being developed in a number of key recruiting areas including Recruiting Strategy, Negotiation Skills, Interviewing, Compliance, Diversity, Benefits, Compensation and advanced applicant tracking system courses.