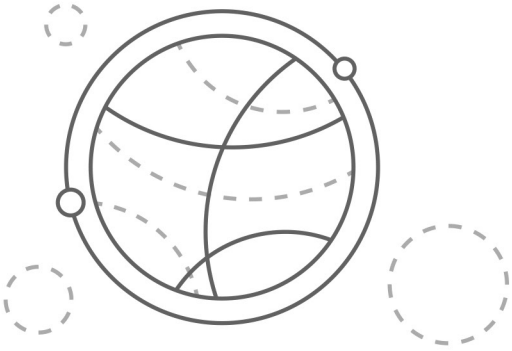

Contact Center Management
on Fast Forward

Succeeding in the New Era of Customer Experience

Brad Cleveland

Fourth Edition

CHAPTER 9:

How Contact Centers Behave

"Every why hath a wherefore."
SHAKESPEARE,
THE COMEDY OF ERRORS

As surely as the laws of physics define the realities of air travel, fundamental principles govern contact centers. Misunderstanding or ignoring these principles leads to inconsistent service, demoralized employees, excessive costs, and poor customer experiences.

Six "immutable laws" are at work in any center that handles inbound contacts. They are immutable because they *always have been* and *always will be* true. All are driven or at least influenced by random contact arrival (see Chapter 3).

Let's take a look at each. Along the way, I'll offer some tips on bringing your team up to speed and ensuring that these immutable laws are working *for* (not *against*) your efforts.

1. When Service Level Goes Up, Occupancy Goes Down

As discussed in Chapter 4, service level is expressed as "X percent of contacts answered in Y seconds." Occupancy is the percentage of time

SIX IMMUTABLE LAWS AT WORK IN CONTACT CENTERS

1. For a given workload, when service level goes up, occupancy goes down.
2. Keep improving service level and you will reach a point of diminishing returns.
3. For a given service level, larger agent groups are more efficient than smaller groups.
4. All other things being equal, pooled groups are more efficient than specialized groups.
5. For a given workload, add staff and average speed of answer will go down.
6. For a given workload, add staff and trunk load will go down.

during an interval that agents who are signed in and handling the workload are actually handling contacts. The inverse of occupancy is the time that agents are waiting to handle contacts.

Average handling time (sec.) = 210; Contacts: 250

Agents	ASA	SL% in 20 Sec.	Occupancy	Trunk Load (in hours)
30	209	24%	97%	54.0
31	75	45%	94%	35.4
32	38	61%	91%	30.2
33	21	73%	88%	28.0
34	13	82%	86%	26.8
35	8	88%	83%	26.1
36	5	92%	81%	25.7
37	3	95%	79%	25.4
38	2	97%	77%	25.3
39	1	98%	75%	25.2
40	1	99%	73%	25.1
41	1	99%	71%	25.1
42	0	100%	69%	25.0

As the table illustrates, a service level of 80 percent of contacts answered in 20 seconds (82/20 in this scenario) equates to an occupancy of 86 percent for that workload. If service level drops to 24 percent answered in 20 seconds, occupancy goes up to 97 percent.

The relationship between occupancy and service level is often misunderstood. The incorrect logic goes something like, "If agents really dig in, service level will go up and so will their occupancy." In reality, if occupancy is high, it is because agents are working on one contact after another, with little or no wait in between. Contacts are stacked up in queue and service level is low. In the worst scenario, occupancy is 100 percent because all customers spend at least some time in queue and agents have no breathing room between contacts.

When service level goes up, occupancy goes down (see figure, Contacts Per Agents Versus Service Level). Therefore, the average contacts handled per agent will also go down. Setting standards on number of contacts handled is not recommended, because agents can't directly control occupancy. Doing so would also conflict with an important objective: ensure that enough agents are available to handle contacts so that your service level objectives are achieved. (We will discuss individual performance standards in Chapter 14.)

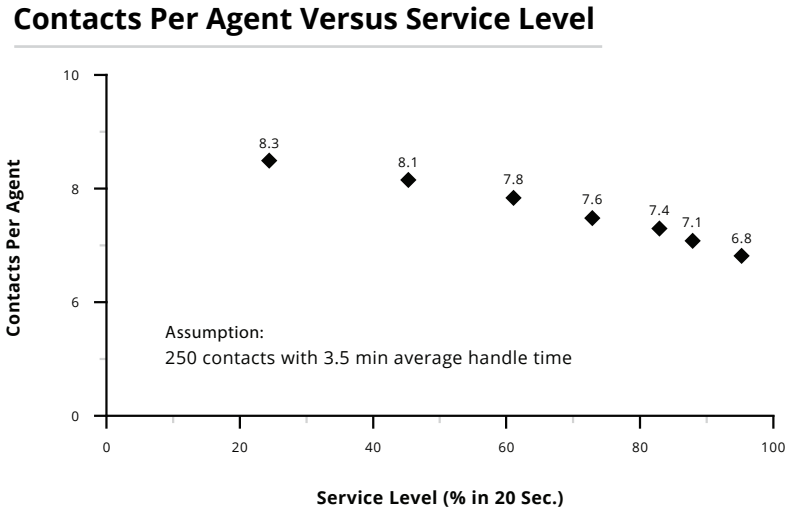

Occupancy is driven by random contact arrival and is heavily influenced by service level and group size (see the third immutable law). Managers don't exactly love this principle—no one wants "unproductive" time baked into the process. However, the time that agents spend waiting for contacts is sliced into 12 seconds here, two seconds there, and so on, and is a factor of how contacts are arriving.

In many centers, agents handle other types of work when the inbound workload slows down. Blended environments (where agents handle different types of work based on workload requirements) make sense—no one has a perfect forecast, so schedules don't always match staff to the call load. (As a rule, successful blended environments don't switch agents from one type of work to another minute-by-minute; that's simply too hard for any person to juggle and quite honestly inefficient. These blended environments do, however, make the switch for larger blocks of time: a few hours of this and a few hours of that.)

But understand what is really happening. When other work is getting done, either: a) you have more agents scheduled than necessary to handle the workload at your service level goal; or b) service level is being sacrificed. Don't try to force occupancy higher through non-contact work than what base staff calculations predict it should be.

When Is Occupancy Too High?

As any agent knows, periods of high occupancy are stressful. Studies suggest that agents begin to burn out around 90 percent occupancy if the condition lasts for an extended time, such as several half hours in a row (some studies set the threshold as low as 88 percent, others as high as 92 percent). Taking breaks is a natural reaction to high occupancy, but it compounds the problem.

Consider this scenario. Jen, Ben, and Mary are three of 32 agents plugged in and taking calls. Staffing calculations (see table above) predict the average occupancy for the half hour for 32 agents will be 91 percent and service

level will be just above 60 percent answered in 20 seconds. This is what these three agents may be thinking:

Jen: *Whew! It's contact after contact this morning. I need a breather. I don't have a scheduled break for 45 minutes, so I'm just going to grab some water for a couple of minutes.*

OOPS. Now there are only 31 agents handling the work. If traffic keeps arriving at the same clip, service level will drop and occupancy will go up.

Ben: *Things are busy today, one contact after another. This customer sure is friendly. I wonder if she's getting that storm I've been hearing about …*

So Ben takes longer on the contact, essentially taking a breather during handling time. Service level drops another notch, and occupancy increases. Mary really begins to feel the load …

Mary: *This contact doesn't really require wrap-up, but …*

This is the start of a downward spiral. If contacts are chronically backed up, service level will consistently be low and occupancy will be high. The real fix, of course, goes to the fundamentals of managing a contact center—a good forecast, accurate staffing calculations, and schedules that match people to the workload. It's also important that each individual understands their impact (see The Power of One).

Occupancy and Adherence to Schedule

Notice an important distinction that this law reveals. When adherence to schedule improves, occupancy goes down. Why? Because when agents are available to handle more contacts, service level will go up. And when service level goes up, occupancy goes down. This means that if your agents adhere to their schedules, they don't have to work as hard. This is an important concept for everyone to understand.

> *When adherence to schedule improves (goes up), occupancy goes down.*

The terms adherence to schedule and occupancy are often incorrectly used interchangeably. They not only mean different things, they move in *opposite* directions. And as we will discuss at in Chapter 14, while adherence to schedule is within the control of individuals, occupancy is determined by factors outside of an individual's control.

2. The Law of Diminishing Returns

Economists identified the law of diminishing returns as it applies to manufacturing many years ago, but it also can have significant impact on other environments, including contact centers. It can be defined this way: when successive individual agents are assigned to a given workload, the incremental improvements in service level that can be attributed to each additional agent will eventually decline.

The Law of Diminishing Returns figure is based on Erlang C data from the

Law of Diminishing Returns

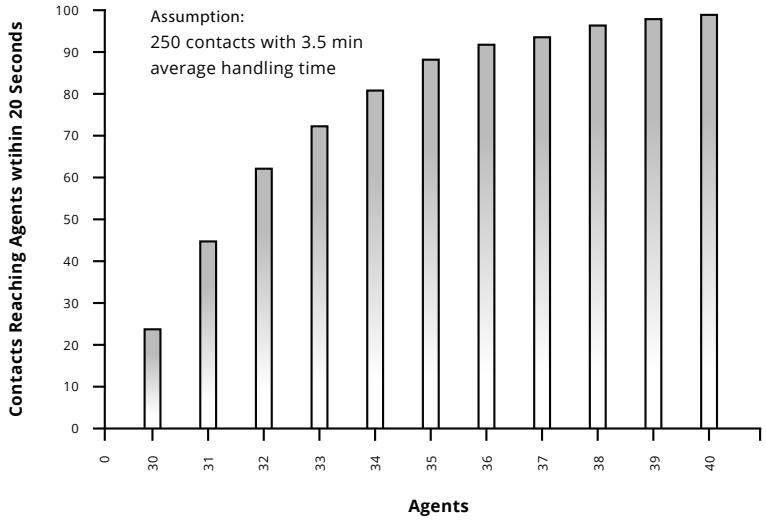

staffing table at the beginning of this chapter. It shows that 30 agents at the given call load will provide a service level of 24 percent in 20 seconds. Keep in mind, these numbers will not be exact—at that low of a service level, many of the contacts will abandon (not get answered at all). But the exact results aside, service level will be poor.

With 31 agents, things improve dramatically, as service level jumps to 45 percent. Adding one more person yields another big improvement. In fact, adding only four or five people takes service level from the depths to something respectable. That means an associated drop in average speed of answer (ASA) and trunk load.

The same principle is true for larger groups, as the next table shows. Each person has a significant positive impact on the queue when service level is low, even in groups that are much larger.

Average handling time (sec.) = 210; Contacts: 1,000 in 1/2 hr.

Agents	ASA	SL% in 20 Sec.	Occupancy	Trunk Load (in hours)
117	607	7%	100%	437.0
118	135	24%	99%	175.0
119	69	39%	98%	138.0
120	42	51%	97%	123.5
121	29	61%	96%	115.9
122	21	69%	96%	111.4
123	15	75%	95%	108.4
124	11	80%	94%	106.3
125	9	85%	93%	104.8
126	7	88%	93%	103.7
127	5	91%	92%	102.9
128	4	93%	91%	102.2
129	3	94%	90%	101.7
130	3	96%	90%	101.4
131	2	97%	89%	101.1
132	2	97%	88%	100.8
133	1	98%	88%	100.7

Contact center managers who struggle with a low service level are fond of this immutable law because it often doesn't take many people to improve things significantly. Those managers who want to be the "best of the best" in terms of service level find that the relationship between varying levels of resources and service level must be clearly outlined in the budgeting process.

Viewed from a different angle, if you have the right number of people handling contacts to begin with, but just a few of them unplug or go unavailable at an inopportune moment, contacts begin to back up. Think of what a stalled car blocking just one lane can quickly do to a busy expressway.

The Power of One

The *power of one* is among the most important principles to introduce to new hires and reinforce with experienced agents. While we know the impact each agent has on individual customers and the subsequent publicity (good or bad) that can come from those experiences, the power of one refers more specifically to queues and wait times.

The central theme that shapes contact center operations is that they are dynamic; workloads arrive randomly in any center that handles customer-initiated contacts. This, coupled with the reality of how queues behave, means that agents who are helping manage the workload affect service level—in a good way—far more than they may realize.

(The concept of the power of one has been used frequently across the business world, in charitable fundraising, and in the contact center profession, including as the title of an excellent booklet by author Penny Reynolds. However, it was first popularized by Australian author Bryce Courtenay, who used it as the title of his 1989 book about a young boy growing up in South Africa.)

Take a look again at the impact of different staffing levels (this is the same table shown at the beginning of the chapter).

Average handling time (sec.) = 210; Contacts: 250

Agents	ASA	SL% in 20 Sec.	Occupancy	Trunk Load (in hours)
30	209	24%	97%	54.0
31	75	45%	94%	35.4
32	38	61%	91%	30.2
33	21	73%	88%	28.0
34	13	82%	86%	26.8
35	8	88%	83%	26.1
36	5	92%	81%	25.7
37	3	95%	79%	25.4
38	2	97%	77%	25.3
39	1	98%	75%	25.2
40	1	99%	73%	25.1
41	1	99%	71%	25.1
42	0	100%	69%	25.0

Service level is bad with 30 agents—just 24 percent answered in 20 seconds. With one additional agent, things improve dramatically. Service level jumps to 45 percent (still not great, but almost twice as good). Average speed of answer drops from 209 to 75 seconds. Occupancy goes down, from 97 percent to 94 percent (that might not sound like a big drop, but it feels a lot better!).

Yes, one person makes *that much* of a difference for customers and the rest of the team! Adding one more person yields another big improvement. As you can see, if you let your eyes follow the rows down the table, there's a point at which adding agents doesn't help much, because service is already good. You get into the law of diminishing returns.

Next, look at what happens to customers at different staffing levels (see figure, Customer Delay). If you have 34 agents handling contacts, 65 customers are waiting five seconds or longer. Seven customers reach agents in the next five seconds, so 58 are still waiting ten seconds or longer. Another

six customers reach agents in the next five seconds, leaving 52 waiting 15 seconds or longer, and so forth. There's still one customer waiting 180 seconds, and no customer waits more than four minutes. It's a very different story, however, if there are only 30 agents handling calls. Dozens of customers are waiting four minutes or longer. The results look far better when just one additional agent is added.

Erlang C for Contact Centers — Customer Delay

Average talk time in seconds: 180 **Average after-contact work in seconds: 30**
Contacts per half hour: 250 **Service level in seconds: 20**

← Number of customers waiting longer than x seconds

Agents	SL%	5	10	15	20	30	40	50	60	90	120	180	240
30	24	203	199	195	191	184	177	170	163	145	129	101	80
31	45	156	149	143	137	126	115	105	97	74	57	34	20
32	61	118	111	104	97	85	74	65	56	38	25	11	5
33	73	89	81	74	67	56	47	39	32	19	11	4	1
34	82	65	58	52	46	37	29	23	18	9	5	1	0
35	88	47	41	36	31	24	18	14	10	4	2	0	0
36	92	34	29	24	21	15	11	8	6	2	1	0	0
37	95	24	20	16	14	9	6	4	3	1	0	0	0
38	97	16	13	11	9	6	4	2	2	0	0	0	0
39	98	11	9	7	5	3	2	1	1	0	0	0	0
40	99	7	6	4	3	2	1	1	0	0	0	0	0
41	99	5	4	3	2	1	1	0	0	0	0	0	0
42	100	3	2	2	1	1	0	0	0	0	0	0	0

Source: ICMI's QueueView Staffing Calculator

Just remember, when queues back up, everybody makes a big difference. Each person has a significant positive impact on customer wait times—which goes far beyond the customers they serve directly. Knowing about these dynamics helps agents understand why schedules are a big deal and why schedule adherence matters.

Experienced customer service employees will correctly point out that the power of one also has a qualitative aspect. Just consider the ripple effect of customer reviews or publicity (good or bad) that can come from any interaction. An American Express study found that consumers tell 21 others on average about a poor service experience. We've all seen videos of bad experiences that went viral. And positive word of mouth builds powerful brands. When you give a good customer service experience, you're creating a powerful marketing force for your company.

The power of one principle is as important as ever, given ever-expanding contact channels and heightened customer expectations for quick and easy service. My encouragement is to keep it front and center with your team!

Here are some of the steps organizations are taking to reinforce the power of one. Think through how you could approach them with your team.

- Educate each person on how much impact he or she has on the queue—and incorporate these or similar scenarios into training.
- Develop reasonable expectations for adherence to schedules, and explain the reasoning behind those expectations.
- Educate everyone on the core steps involved in forecasting and resource planning, so that they know how schedules are produced and where they come from.
- Provide real-time queue information to agents on readerboards, desktops or phones. (We'll discuss how to use this information in Chapter 11.)
- Develop appropriate priorities for the full range of tasks that your team handles and guidelines for how to respond to evolving conditions.

UNDERSTANDING HOW CONTACT CENTERS BEHAVE

The central reality that shapes contact center operations is that they are dynamic in the truest sense. Because of the randomly arriving nature of customer contacts, each person has a big impact on the organization's responsiveness. That, in turn, is an important enabler to delivering great customer experiences, boosting loyalty and contributing to successful business results.

A notable trend among the most effective contact centers is to educate their entire teams (agents, supervisors, managers and analysts, as well as colleagues from across their organizations) on contact center dynamics, the power of one, and the value the organization delivers when strong cross-functional support is in place. American Express, USAA, and FedEx are just a few examples of highly rated companies that have made this an ongoing priority.

Providing an understanding of how contact centers behave is a gift to those who work in or support them. It just makes good sense.

3. Larger Groups Are More Efficient

Average group productivity (contacts that a group handles) is not a constant factor. Instead, it's constantly fluctuating as the workload ebbs and flows. Even when you maintain a consistent service level through good planning and on-target scheduling; you'll find that average productivity is relatively lower at lower volumes and relatively higher at higher workloads. Because the number of contacts is changing throughout the day, so is average group productivity.

Why? Mathematically, larger groups of agents are more efficient than smaller groups, at the same service level. Therefore, larger groups assigned to heavy mid-morning traffic will be more efficient than smaller groups handling the lighter evening load. So, calculating staff the wrong way—assuming fixed productivity—will be highly inaccurate (see table, The Impact of Group Size).

The Impact of Group Size

Contacts in Half Hour	Service Level	Agents Required	Occupancy	Avg. Contacts Per Agent
50	80/20	9	65%	5.6
100	80/20	15	78%	6.7
500	80/20	65	90%	7.7
1,000	80/20	124	94%	8.1

Assumption: Average Handle Time is 3.5 minutes

This is yet another reason why setting standards on the number of contacts that agents handle is an unfair way to measure productivity. Attempting to compare groups or sites in a multi-site environment may also be misleading (the exception would be a network that finds the longest-waiting agent, a true virtual contact center).

Despite mathematical efficiencies, there is a point where groups become so large that occupancy becomes too high for agents. Some managers believe that the number of agents in a single group should be limited to 125 to

Fallacy of Attempting to Achieve a Constant Service Level With Fixed Productivity

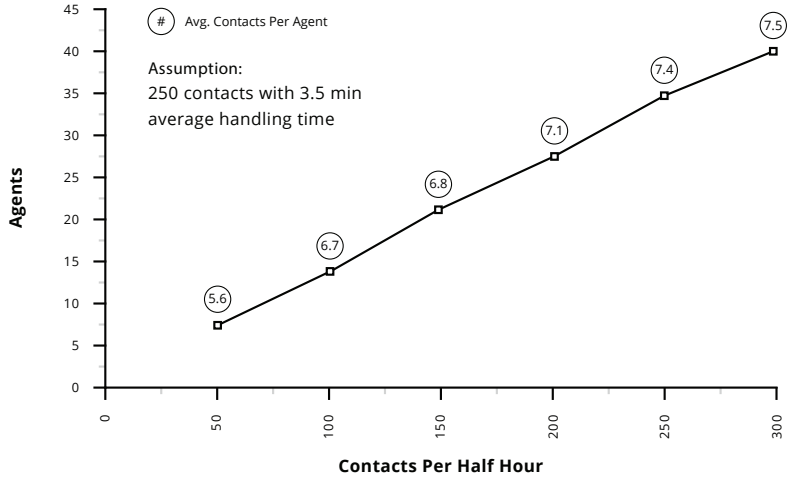

150 people. However, plenty of centers have much larger agent groups (the U.S. Social Security Administration, Centrelink Australia, China Mobile and others have hundreds or even thousands of agents in a single agent group).

Rather than establishing a strict limit to group size, a better approach is to watch occupancy and take appropriate measures when it climbs above 90 percent. For example, in the scenario on page 207, scheduling so that 129 agents or more are handling the work is recommended, even though the required service level may be exceeded. Customers sure won't mind, and your staff will be able to function efficiently throughout their shift.

4. The Powerful Pooling Principle

The powerful pooling principle is a mathematical fact that goes like this: any movement in the direction of consolidation of resources will result in improved traffic-carrying efficiency. Conversely, any movement away from consolidation of resources will result in reduced traffic-carrying efficiency. Put more simply, if you take several small, specialized agent groups, and effectively cross-train them and put them into a single group, you'll have a more efficient environment.

THE POWERFUL POOLING PRINCIPLE

- Handle more contacts, at the same service level, with the same number of agents
- Handle the same number of contacts, at the same service level, with fewer agents
- Handle the same number of contacts, at a better service level, with the same number of agents

Note, again, the table on page 213, which compares service level to group size. Fifteen agents are required to provide a service level of 80/20. But only 124 agents, not 150, are necessary to handle a load 10 times as large.

The pooling principle should be a consideration from the highest levels of strategic planning (How many centers should you have? How should agent groups be designed?), down to more tactical decisions related to real-time adjustments or how to best invest training time and resources.

In one sense, pooling resources is at the heart of what ACDs and networks do. A clear trend in recent years, though, is the recognition that customers often have different needs and expectations, and that different agents with a mix of aptitudes and skills are required. Powerful capabilities, such as skills-based routing, give us the means to route and handle contacts based on myriad criteria (see Chapter 7).

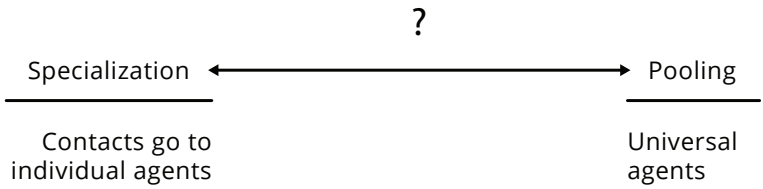

Can you have specialization without forgoing the benefits of the powerful pooling principle? It depends. Skills-based routing can boost efficiency by getting contacts to the agents best suited to handle them. But when not managed well, or when overused, the number of contingencies can multiply beyond your management team's ability to understand and manage them. The interplay can become stupefying. And the whole notion of agent groups and pooling begins to erode.

When skills and routing priorities become too complex, related dangers begin to emerge. Doug Casterton, Head of Global Workforce Planning and Scheduling for Trip Advisor in Oxford, U.K., warns of the "ever-sinking queue." When skill priorities have been set a different level, "it's possible for the higher priority contacts to jump the queue, and if you have not correctly staffed, the lower priority contacts may never actually reach an agent."

As real and pervasive as the pooling principle is, it is not an all-or-nothing proposition. There is a continuum between pooling and specialization—think of a variable thermostat rather than an on/off switch. Your objective should be to get as close to the pooled end of the spectrum as circumstances allow. Examples of supporting steps would include:

- Improve training and coaching, to enable agents to handle more contact types.
- Hire multilingual agents to better cover all supported languages.
- Integrate new channels—such as social media or chat—into existing agent groups as much as feasible.
- Improve knowledge management systems so that agents have the information needed to handle a broad a range of contacts.
- Work on reducing turnover and improving agent tenure (their experience levels).

You get the idea. These and other steps you can take to effectively broaden the work types agents handle will, by definition, boost efficiencies.

5. Add Staff and ASA Goes Down

Anyone who has ever waited in line for anything knows that if there were a few more tollbooths, open check-out aisles, or people behind the counter, the line wouldn't be so long! And when someone behind the counter gets reassigned to another task or goes on break, the wait increases (this happens anytime I enter a physical line—perhaps it's a cosmic joke on those of us who study queues).

The same principle is at work in contact centers. When more agents are plugged in and handling contacts, assuming they are trained to do so proficiently, the queue will be shorter. Fewer agents means a longer queue. This principle leads to the next immutable law.

6. Add Staff and Trunk Load Goes Down

When more agents are assigned to a given workload, trunk load (the load on the network that handles voice and data) goes down. The converse is also true: when fewer agents are available to handle a given workload, trunk load goes up because the delay increases (see discussion on trunks, Chapter 7).

Each customer connected to your system is part of the workload, whether they are talking to an agent or waiting in queue. If you have toll-free service (or any other service that charges a usage fee), you are paying for this time. Telecommunications costs are inextricably wrapped in staffing issues. If service level is continually low, the costs of network services will escalate.

The following example illustrates the tradeoffs between staffing levels and service level, average speed of answer, occupancy, and trunk load. Recall from Chapter 7 that trunk load represents how much time (in hours) customers are queued for and/or talking to agents in a particular group over the equivalent of an hour. Staff is calculated for a half hour's traffic, but the trunk load is converted to an hour's traffic ("erlangs") simply because telecom managers universally use hour increments for engineering and management purposes.

Average handling time (sec.)* = 210; Contacts: 350 in half hour

Agents	ASA	SL% in 20 Sec.	Occupancy	Trunk Load (in hours)
42	144	29%	97%	62.9
43	63	58%	95%	47.3
44	35	61%	93%	41.8
45	21	72%	91%	39.1
46	13	80%	89%	37.6
47	9	86%	87%	36.7
48	6	90%	85%	36.1
49	4	93%	83%	35.8
50	3	95%	82%	35.5
51	2	97%	80%	35.3
52	1	98%	79%	35.2

* Avg Talk Time 180 sec., Avg After Call Work 30 sec.

Understanding Trunk Load

Want to see where these numbers come from? (Not to worry; it doesn't take long, and you'll get the idea). Using the scenario in the table, assume that you have 46 agents handling contacts and, therefore, will be able to achieve a service level of 80/20. Here's how the calculations produced an estimated 37.6 hours on the trunks:

- First, you can see that customers will be queued for agents an average of 13 seconds (ASA) and will be connected to agents an average of 180 seconds (average talk time), for a total of 193 seconds. The 180 seconds represents the forecast for what average talk time will likely be; the 13 seconds ASA comes from the Erlang C calculation.

- Because the table provides volume for a half hour, double 350 contacts to assume 700 contacts in an hour.

- Because the 700 calls spend an average 193 seconds queuing for and connected to agents, the trunk load in seconds is 135,100 seconds (700 contacts × 193 seconds).

- Finally, because trunk load is presented in erlangs (hours of traffic over the course of an hour), divide 135,100 by 3,600 (the number of seconds in an hour) and you come up with 37.6 hours. To use the correct telecom lingo, you'll have the equivalent of 37.6 erlangs of traffic on the trunks for this agent group during this time period. (Note: this example does not include the time customers may spend in the IVR before arriving at the agent group, which would need to be added.)

The big variable is *the time customers spend in queue before they get connected*. It goes up (gets worse) with fewer agents, and goes down (gets better) with more agents. Glance through the table, and you'll see that if 50 agents are handling calls, ASA will be a projected 3 seconds. If 42 agents are handling calls, ASA will be a projected 144 seconds. In short, the number of agents handling work determines the average delay, which is a key variable in trunk load and, accordingly, in what you pay for toll-free services.

The Impact of Staff on Toll-Free Costs

As you decide on service levels and how to allocate budgets, you should think about network costs, too. All other things being equal, if your service level is low, adding an agent will often bring total costs down because network costs will drop dramatically.

The staff versus toll-free costs tradeoff used to be much more dramatic. In many parts of the world, toll-free costs are just pennies (or less) a minute. Toll-free service used to be much more expensive to organizations (15 cents per minute or higher). So, improving service level meant huge drops in network costs, often producing savings that far surpassed the cost of adding staff. With today's far lower network costs, the tradeoff is less significant.

However, you still have expense for carrying the traffic, and there are also costs related to ports, IVR capacity, maintenance, taxes, and other budgetary line items. Delay takes resources and it is not free. Assessing the impact on network costs underscores the importance of considering both agent and network costs together. Improving service level will save money on network services; these savings should be factored into predictions of overall costs.

The Cost of Delay

The direct expense of putting customers in queue is called the "cost of delay." It is expressed in terms of how much you pay for toll-free service each day (or month, hour, or half hour) just for customers to wait in queue until they reach an agent.

You may want to plot the cost of delay. It's simple. First, take the total delay for the day, as reported by your ACD, and convert that into minutes or hours. Next, multiply the minutes or hours of delay by the average per-minute or per-hour cost of your toll-free service. Then add that figure to a graph that illustrates these costs (see example).

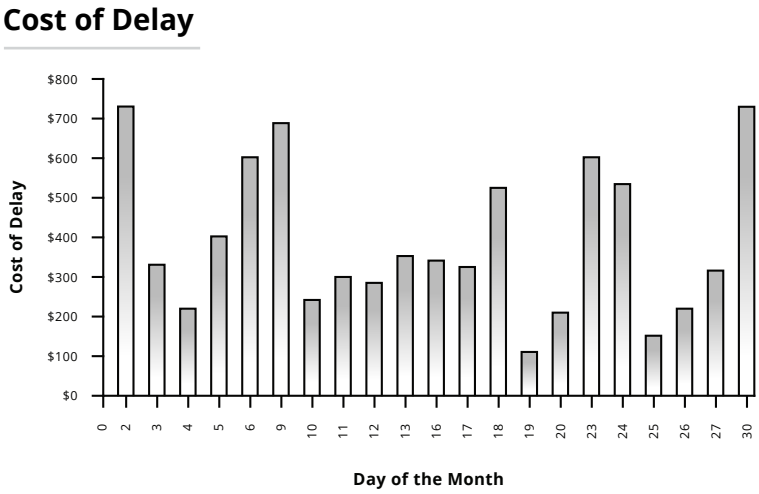

This graph will be a reminder that poor service is not cheap. And it will catch the interest of senior managers, who will look at the graph and wonder aloud, "You mean that's what we're paying just for customers to wait? Why, we could use that money for ..."

Points to Remember

- There are immutable laws at work in any center that handles inbound contacts.
- A common theme runs through these laws: do a good job of matching staff with the workload or both customers and agents will suffer the consequences.
- The burden doesn't fall solely on those who do the planning and scheduling. Designing and managing a contact center requires a big-picture perspective and the collaborative effort of all involved.
- A good understanding of these immutable laws is necessary for developing an accurate planning process, setting fair objectives and standards, developing a good strategy, and just about every other aspect of effective management.

- It's important that agents, senior-level managers, and others who work in or support the contact center are aware of these principles.

Brad Cleveland is known globally as one of today's foremost experts in customer strategy and management. He has worked across 45 U.S. states and in over 60 countries for clients as diverse as American Express, Apple, USAA, the University of California, and the federal governments of Australia, Canada, and the United States. His books and articles have been translated into more than a dozen languages. He is a sought-after consultant and a popular speaker who presents with sparkle, insight, and humor.

Connect with Brad:

Web: www.BradCleveland.com
Twitter: @bradcleveland
Email: info@bradcleveland.com

ICMI™

ICMI has been at the forefront of improving customer experiences for over 30 years. We are 100% focused on empowering contact center professionals to improve their skills, advance their career, and enhance the performance of their organizations. More than 50,000 organizations worldwide have joined ICMI on a collective journey to elevate the customer experience through training, events, consulting, and information resources.

Whether it be measuring a contact center's effectiveness, identifying their omnichannel strategy, assessing their operations, providing tools to manage and motivate their teams, or creating a technology roadmap, ICMI's experienced and dedicated team of industry insiders, trainers, speakers and consultants are committed to helping you raise the strategic value of your contact center. Learn more at ICMI.com.