MetricNet
Performance Benchmarking

MetricNet at FUSION 14

Metric of the Month

Each month MetricNet highlights one Key Performance Indicator for the Service Desk or Desktop Support. We define the KPI, provide recent benchmarking data for the metric, and discuss key correlations and cause/effect relationships for the metric. The purpose is to familiarize you with the Key Performance Indicators that really matter to your support organization, and to provide actionable insight on how to leverage these KPI's to improve your performance. This eBook contains all of our Metric of the Month articles in one place!

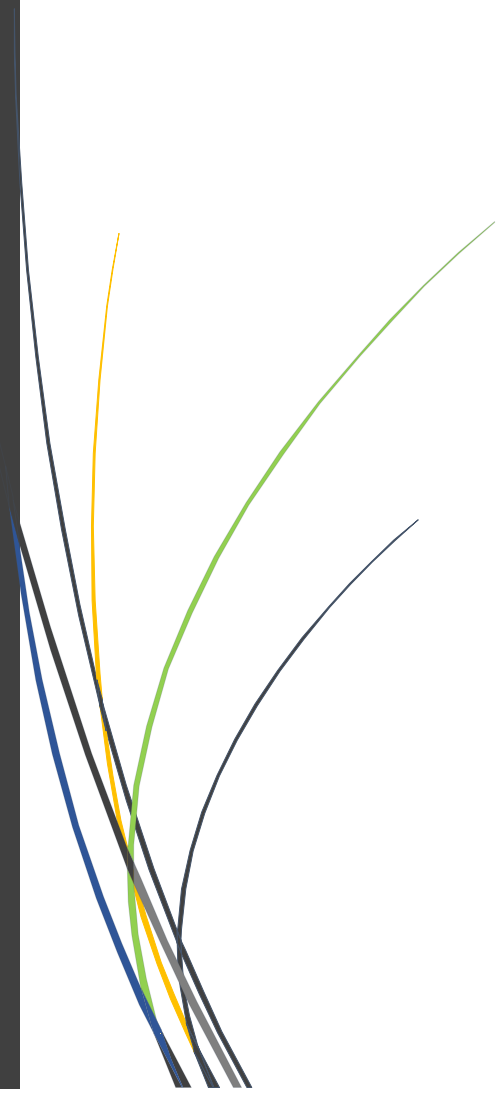

MetricNet Certified eBook
WWW.METRICNET.COM

Contents

First Contact Resolution Rate

Customers tend to be impatient when they want service. It doesn't matter if they are calling their bank, their cable company, or their service desk. They want a resolution to their problem or an answer to their question right then and there! In fact, research across many different industries bears this out. Customer satisfaction – for virtually any type of customer service – is strongly correlated with First Contact Resolution.

For a service desk, First Contact Resolution (FCR) is the percentage of contacts that are resolved by the Service Desk on the first interaction with the customer. For live calls or web chats, this means that the customer's issue is resolved before they hang up the phone, or end the chat session. Calls or chats that require a customer callback, or are escalated to another source of support do not qualify for first contact resolution. For emails, which now account for a significant percentage of all service desk contacts, the de facto standard emerging in the industry is that resolution within one business hour of receiving a customer email counts as FCR.

FCR is typically measured in one of two ways: 1) the service desk agent checks a box on the trouble ticket at the conclusion of the call or chat session to indicate if the call was successfully concluded on the initial contact, or 2) customers are asked in follow-up customer satisfaction surveys whether their call was resolved and concluded on the initial contact with the service desk. The first method requires periodic audits to ensure that agents are accurately reporting FCR on the tickets they handle. This is done by reviewing a representative sample of tickets each month to determine if the tickets designated FCR by an agent are, in fact, being resolved on the first contact with the customer. Neither method of measuring FCR is perfect, but it is one of the most important KPI's to track and trend. Here's why…

Why it's Important

A high FCR is almost always associated with high levels of customer satisfaction. FCR is a measure of how effectively your service desk conducts its business, and is a function of many factors, including the complexity and types of transactions handled, the experience of your agents, the quality of agent training, and tools such as knowledge management and remote diagnostics that are available to your agents.

Figure 1 below shows the relationship between FCR and Customer Satisfaction for a representative cross-section of service desks worldwide. This strong cause-and-effect relationship should come as no surprise. As stated above, we all want closure on the first contact with our service providers.

Figure 1: Net FCR vs. Customer Satisfaction

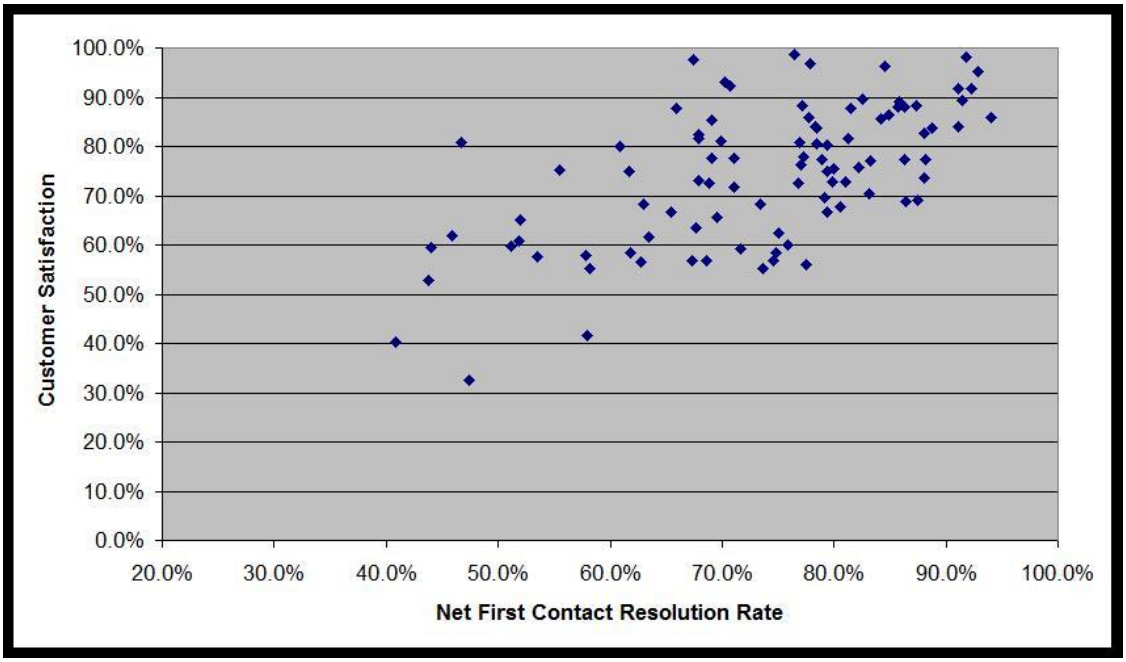

Net vs. Gross FCR

Certain types of contacts cannot be resolved remotely by the service desk. These include hardware break/fix, and physical move/add/change requests. So do these contacts count against a service desk when calculating FCR? Well, yes and no. Let me explain.

Gross FCR looks at all incoming contacts and makes no adjustment for contacts that cannot be resolved by the level 1 service desk. The formula for Gross FCR is:

Gross FCR = (Number of contacts resolved initially ÷ All incoming contacts)

By contrast, *Net FCR* makes adjustments for contacts that cannot be resolved remotely at the service desk. Specifically, the calculation for Net FCR carves out those contacts that cannot be resolved at level 1. The formula for Net FCR is:

Net FCR = (Number of contacts resolved initially ÷ (All incoming contacts – contacts that cannot be resolved at level 1))

In other words, the denominator of the Net FCR ratio is adjusted to include only contacts that can *potentially* be resolved at level 1. These adjustments are sometimes called carve outs.

Net FCR is by far the more relevant of the two FCR metrics. In fact, most organizations do not even track Gross FCR because it produces a distorted picture of the service desk, and is often misinterpreted.

Key Drivers of FCR

As shown above, FCR is a key driver – in fact, the most important driver – of Customer Satisfaction. But what drives FCR? If a service desk wants to improve FCR, how would they go about it? The biggest driver of FCR is Agent Training hours. Figures 2 and 3 below show the

relationship between New Agent Training Hours, Annual Agent Training Hours, and Net First Contact Resolution Rate. The clear implication is that agent training pays off in terms of improved FCR, and that, in turn, yields improvements in Customer Satisfaction.

Figure 2: New Agent Training Hours vs. Net FCR

Figure 3: Annual Agent Training Hours vs. Net FCR

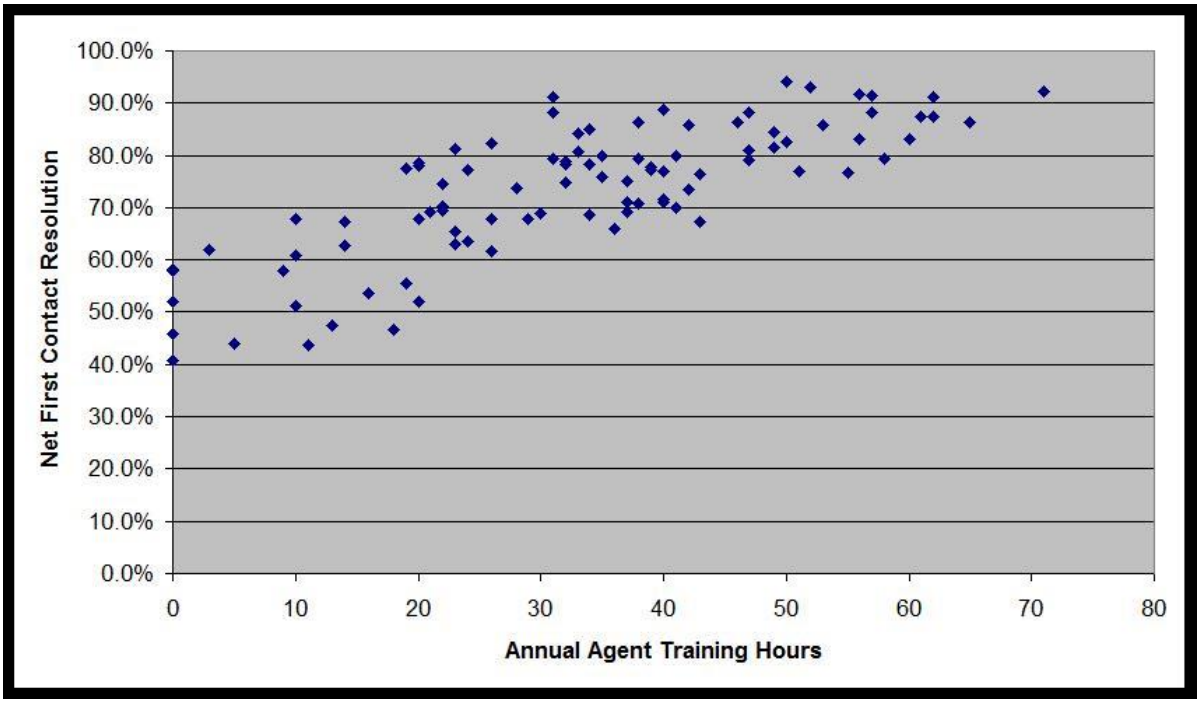

Benchmark Ranges for FCR

MetricNet's benchmarking database shows that the average Net FCR for service desks worldwide is about 74%. This number varies widely, however, from a low of about 41%, to a high of 94%. Those at the low end of this range are predominantly "log-and-dispatch" service desks that resolve simple issues such as password resets, but dispatch anything more complex to another level of support. Those performing at the upper end of this range generally have highly trained agents that are equipped with tools such as knowledge management systems and remote diagnostic capabilities that enable agents to resolve more than 90% of eligible contacts on the first interaction with the customer.

% Resolved Level 1 Capable

% Resolved Level 1 Capable is a Desktop Support metric. It measures the percentage of tickets resolved by desktop support that could have been resolved by the level 1 service desk. This happens when the service desk dispatches or escalates a ticket to desktop support that could have been resolved by the service desk, or when a user bypasses the service desk altogether, and goes directly to desktop support for a resolution to their problem. Although the metric is tracked at desktop support, it has strong implications for both desktop support and the service desk.

Let's say, for example, that a desktop support group handles an average of 1,000 tickets per month. Let's further assume that 200 of these tickets could have been resolved by the service desk. The value for % Resolved Level 1 Capable would be 20% (200 tickets *resolvable* at level 1 ÷ 1,000 tickets *resolved* by desktop support).

Ideally, this metric should be as low as possible because it costs much more to resolve a ticket at the site of a user (desktop support) than it costs to resolve a ticket remotely (by the level 1 service desk). Indeed, "Get it Done at Level 1!" has become a common rallying cry for many support organizations who recognize that one way to minimize support costs is to resolve as many tickets remotely as possible – at the service desk.

% Resolved Level 1 Capable is typically tracked when a desktop support technician closes out a ticket. If the judgment of the technician is that the ticket could have been resolved by the service desk, they will check a box on the ticket indicating that the ticket was resolvable at Level 1. This method requires periodic audits to ensure that desktop support technicians are accurately reporting tickets that could have been resolved by the service desk. An audit can be performed by reviewing a representative sampling of tickets each month to determine if the tickets designated resolvable at level 1 could, in fact, have been resolved by the service desk. Likewise, tickets not designated as resolvable at level 1 are sampled to determine if tickets are routinely being missed that should have been designated as resolvable at level 1.

Why it's Important

Total Cost of Ownership (TCO) for end user support includes the cost of support from all sources, including the service desk, desktop support, other groups or individuals in IT, and vendors. It is a well established fact that the cost of resolving a ticket at level 1 is lower than the cost of resolving a ticket at desktop support. Figure 1 below shows the average cost per ticket for resolution at various support levels in an organization. These figures represent fully loaded costs, including all personnel salaries and benefits, desktop technology and software licensing fees, telecom, facilities, travel, training, and office supplies. Furthermore, these costs are additive: if a user first contacts the service desk, and the service desk in turn opens a ticket and dispatches it to desktop support, the total cost of resolution is $22 for the service desk *plus* an additional $62 for desktop support,

for a total of $84. The clear message here is that getting it done at level 1 is not just a catchy phrase: it can significantly reduce your support costs!

Figure 1: Cost per Ticket – North American Averages

Support Level	Cost per Ticket
Vendor	$471
Field Support	$196
Level 3 IT (apps, networking, NOC, etc.)	$85
Level 2: Desktop Support	$62
Level 1: Service Desk	$22

Minimizing Defects

Every ticket resolved at desktop support that could have been resolved by the level 1 service desk represents a defect in the support organization. You can estimate the cost of these defects by multiplying the number of defective tickets by the cost of resolution at desktop support. Continuing with our example from above with 200 tickets that could have been resolved at level 1, and multiplying that by the $62 per desktop support ticket from Figure 1, we get an estimated defect cost of $12,400 per month. The cost of these escalation defects can really add up! So how do we minimize them?

Here are four suggestions for minimizing the % Resolved Level 1 Capable:

1. Begin tracking the metric at desktop support. You cannot control, let alone reduce the number of ticket defects until you begin tracking % Resolved Level 1 Capable.

2. Conduct root cause analysis on the tickets that are being escalated to desktop support that should have been resolved by the service desk. Are there common themes? Are they coming disproportionately from one or more agents? How many are the result of customers bypassing the service desk, and going directly to desktop support with their support requests?

3. Provide targeted training at level 1 to increase awareness of the importance of this metric, and to reduce the number of tickets dispatched to desktop support that should be resolved by the service desk.

4. Insist upon a strict Single Point of Contact (SPOC) support model, whereby all customer support requests go through the service desk. This will prevent "drive-bys", which happen when desktop support technicians get pulled into an on-the-spot support request without a ticket first being logged and dispatched by the service desk. Figure 2 below shows that support organizations that follow a strict SPOC support model average 15.3% for % Resolved Level 1 Capable, while those that do not follow a strict SPOC model average 22.8%.

Figure 2: % Resolved Level 1 Capable

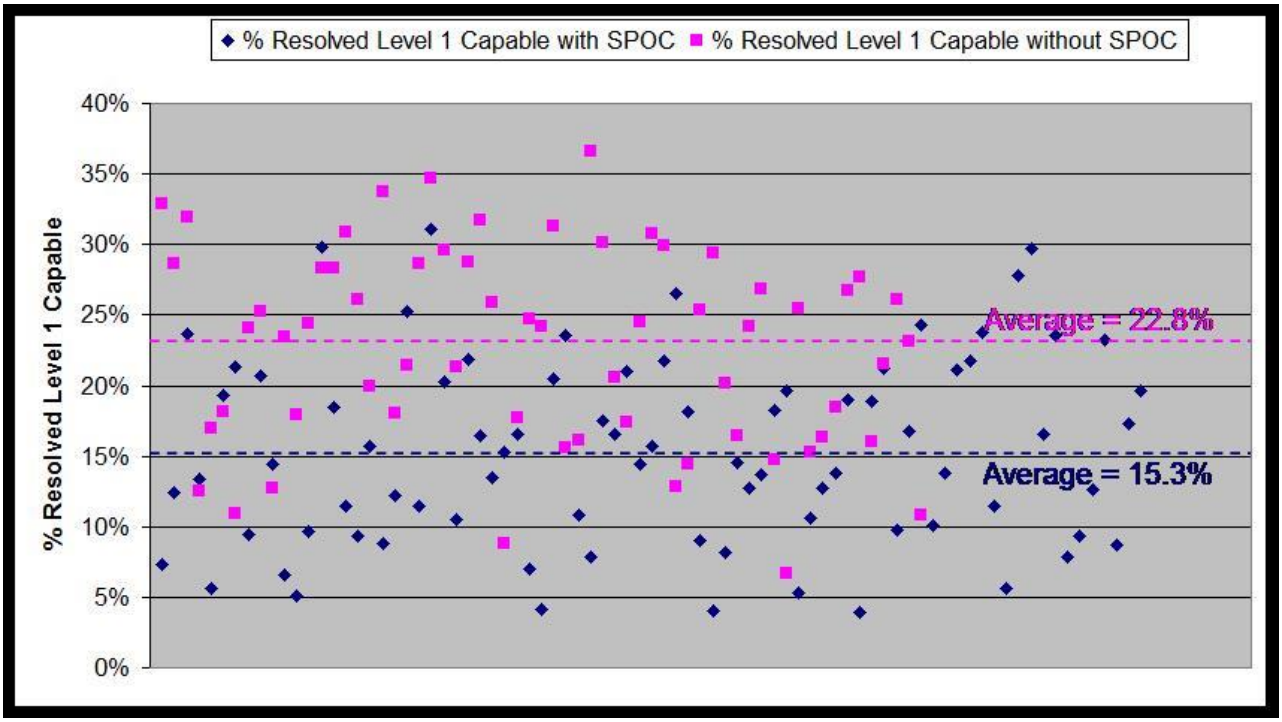

Benchmark Data for % Resolved Level 1 Capable

MetricNet's benchmarking database shows that % Resolved Level 1 Capable ranges from a low of 4%, to a high of 37%, with an average value of 18.6%. That is, 18.6% of all tickets closed by desktop support could have, and should have been resolved by the service desk. This is a surprisingly large number, and indicates that most support organizations have ample opportunity to reduce their support costs by managing this metric more aggressively. To perform in the top quartile for this metric, your % Resolved Level 1 Capable should be less than 12.7%. Figure 3 below shows the global data distribution for % Resolved Level 1 Capable.

Figure 3: Global Data for % Resolved Level 1 Capable

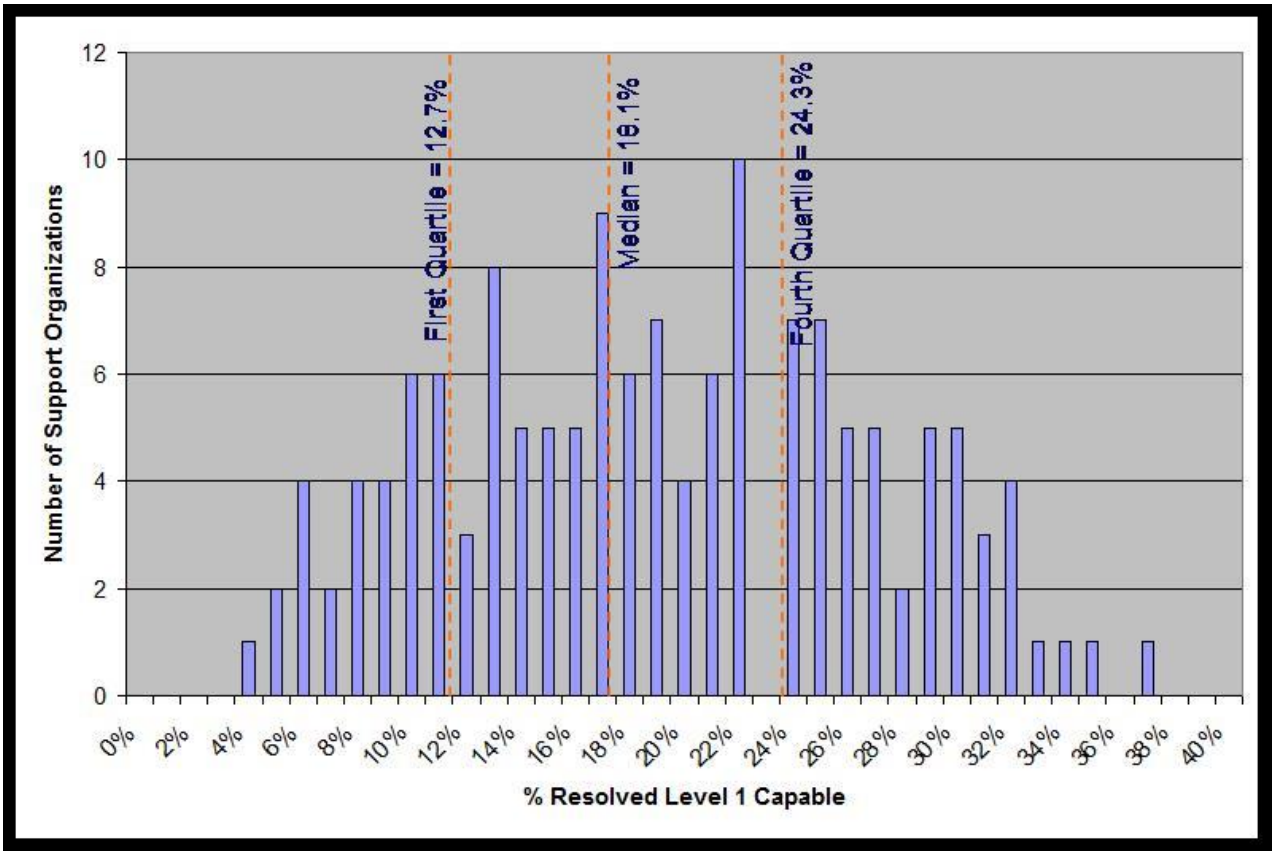
First Quartile = 12.7%
Median = 18.1%
Fourth Quartile = 24.3%
Number of Support Organizations
% Resolved Level 1 Capable
0
2
4
6
8
10
12
0%
2%
4%
6%
8%
10%
12%
14%
16%
18%
20%
22%
24%
26%
28%
30%
32%
34%
36%
38%
40%

Agent Job Satisfaction

Customer Satisfaction is top-of-mind for virtually every service organization. And for good reason: it is the single most important measure of quality for a service desk or desktop support group. But what about Agent Satisfaction? How important is that, and why don't more service desks track this metric? It turns out that it's plenty important, and every support organization should track and trend this metric on an ongoing basis.

Agent Satisfaction is the percentage of agents on the service desk that are either satisfied or very satisfied with their job. It is typically measured annually or semi-annually using an Agent Satisfaction Survey. Unfortunately, fewer than 30% of all service desks track Agent Satisfaction. When I ask clients why they don't track this metric, the answer is usually the same: it's too difficult to measure, or the metric just doesn't matter. On both counts, they are wrong. Here's why…

Why it's Important

Agent Satisfaction is a bellwether metric that impacts many other metrics in the service desk. It is positively correlated with Customer satisfaction, and negatively correlated with Agent Absenteeism and Turnover, meaning that absenteeism and turnover go down as Agent Satisfaction goes up. Figures 1, 2, and 3 below show these correlations using data from MetricNet's 2011 service desk benchmarking database.

Figure 1: Agent Job Satisfaction vs. Customer Satisfaction

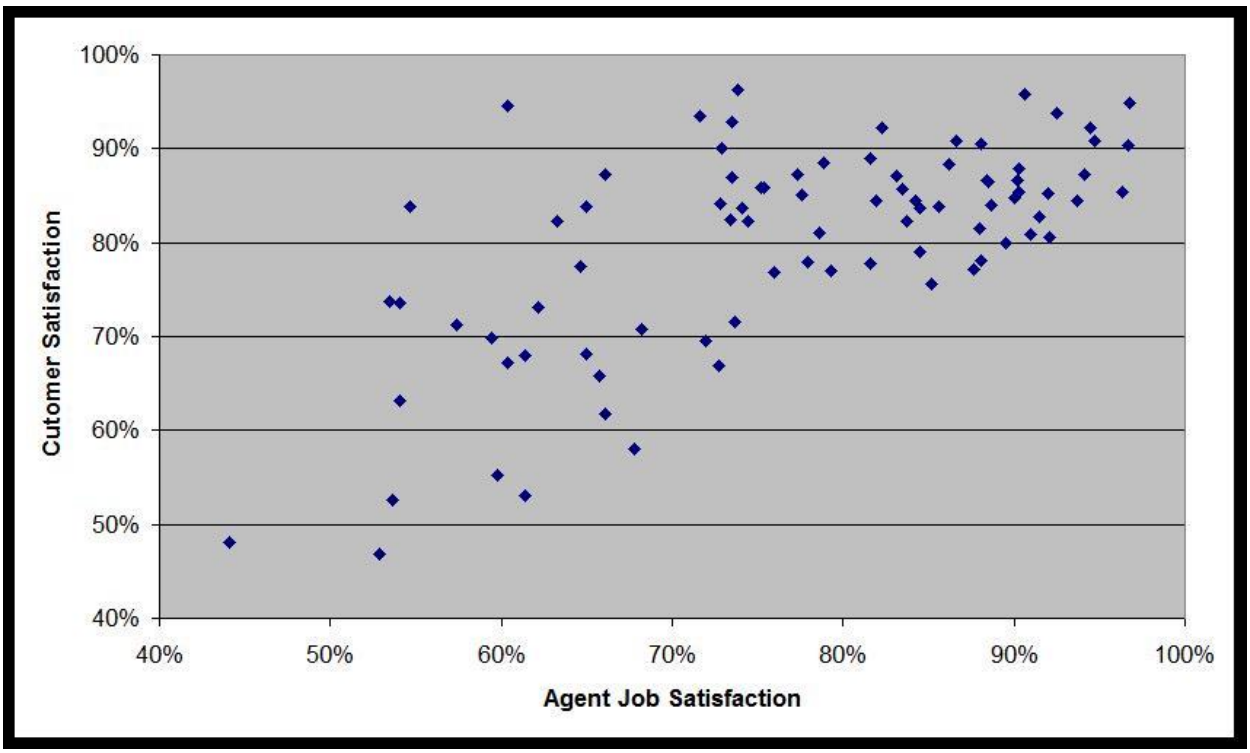

Figure 2: Agent Job Satisfaction vs. Annual Agent Turnover

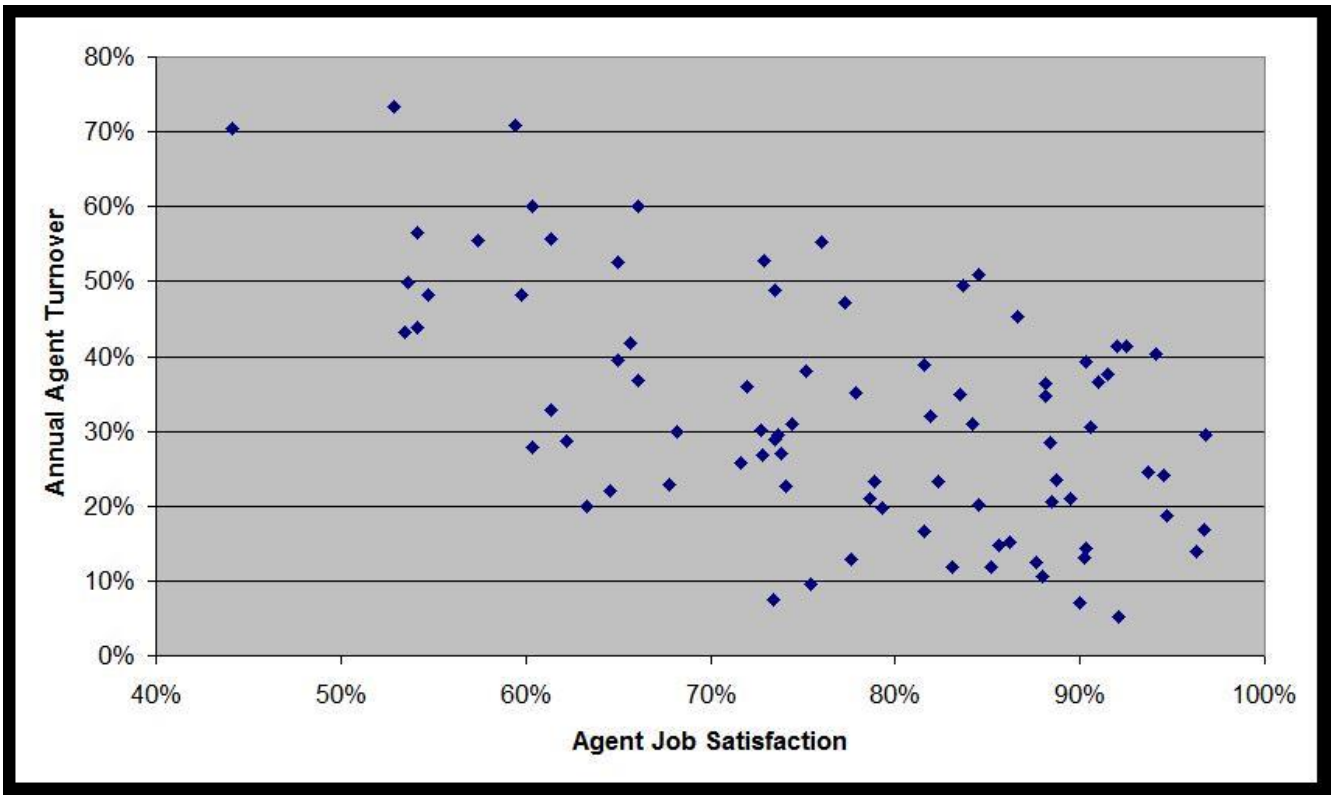

Figure 3: Agent Job Satisfaction vs. Daily Agent Absenteeism

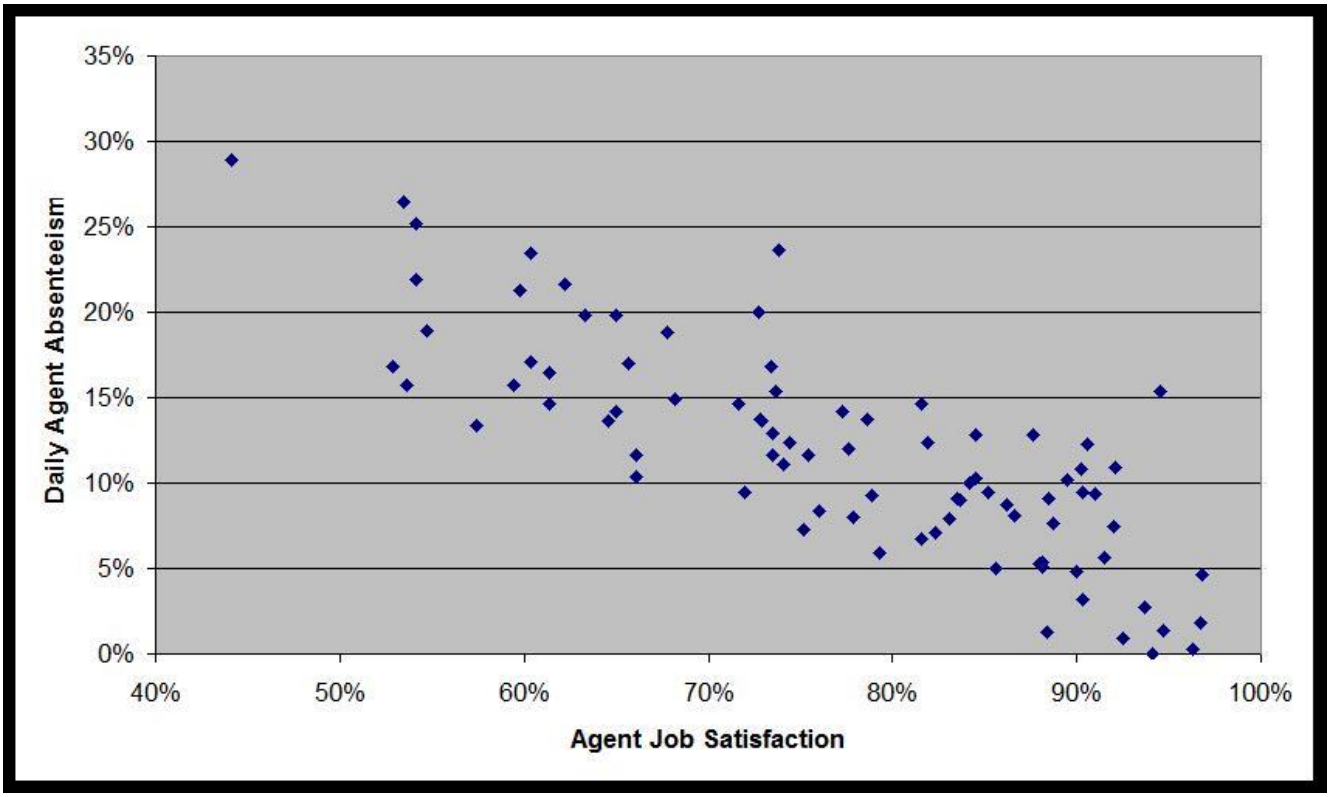

So, why do we care about these correlations? Two reasons. First, they are not just spurious correlations; they are cause-and-effect relationships. Agent Satisfaction (along with FCR) *drives* customer satisfaction; it *drives*agent turnover; and it *drives* agent absenteeism. More importantly, if we can control Agent Satisfaction (which we can), then we can drive positive improvements in customer satisfaction, turnover, and absenteeism.

The second reason is that high Agent Satisfaction has positive economic benefits. Inasmuch as improvements in Agent Satisfaction can reduce agent turnover, it can also reduce the costs associated with turnover. Last year in North America the direct cost of replacing an agent, including the cost of screening, recruiting, interviewing, and training an agent was more than $10,000. And this does not even take into account the value of the knowledge and expertise that is lost when an agent leaves the service desk. Turnover is indeed costly!

Key Drivers of Agent Satisfaction

I mentioned above that Agent Satisfaction can be controlled. But how? Many service desks struggle with low morale among their agents, and appear unable, despite their best efforts, to maintain a high level of morale in their service desks. Can we do something *proactive* to manage and improve agent job satisfaction? The answer is yes, and the solution is surprisingly simple.

It turns out that Agent Satisfaction is driven by training hours, coaching hours, and the existence of a formal career path. Once again we can turn to our benchmarking data to make this point. Figures 4 and 5 below show how training and career pathing impact Agent Job Satisfaction.

Figure 4: Annual Training Hours vs. Agent Job Satisfaction

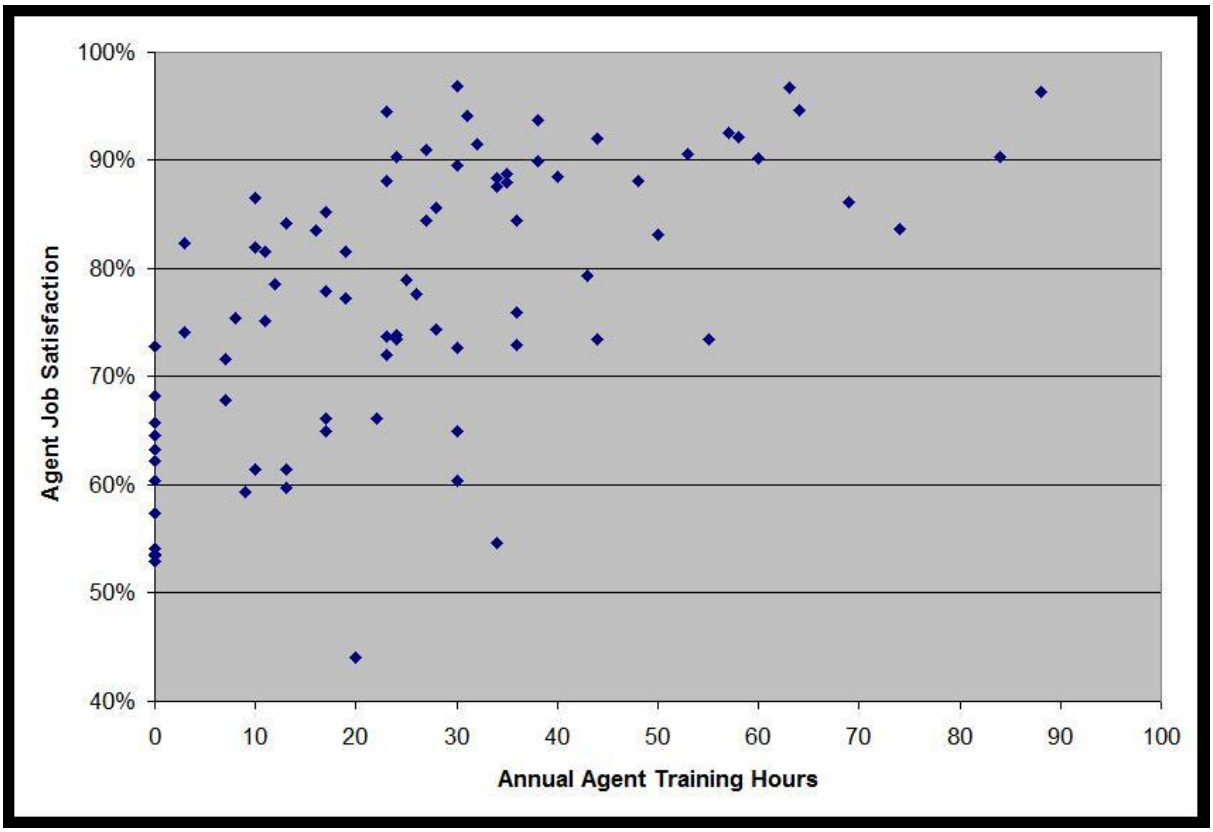

Figure 5: Career Pathing and Agent Job Satisfaction

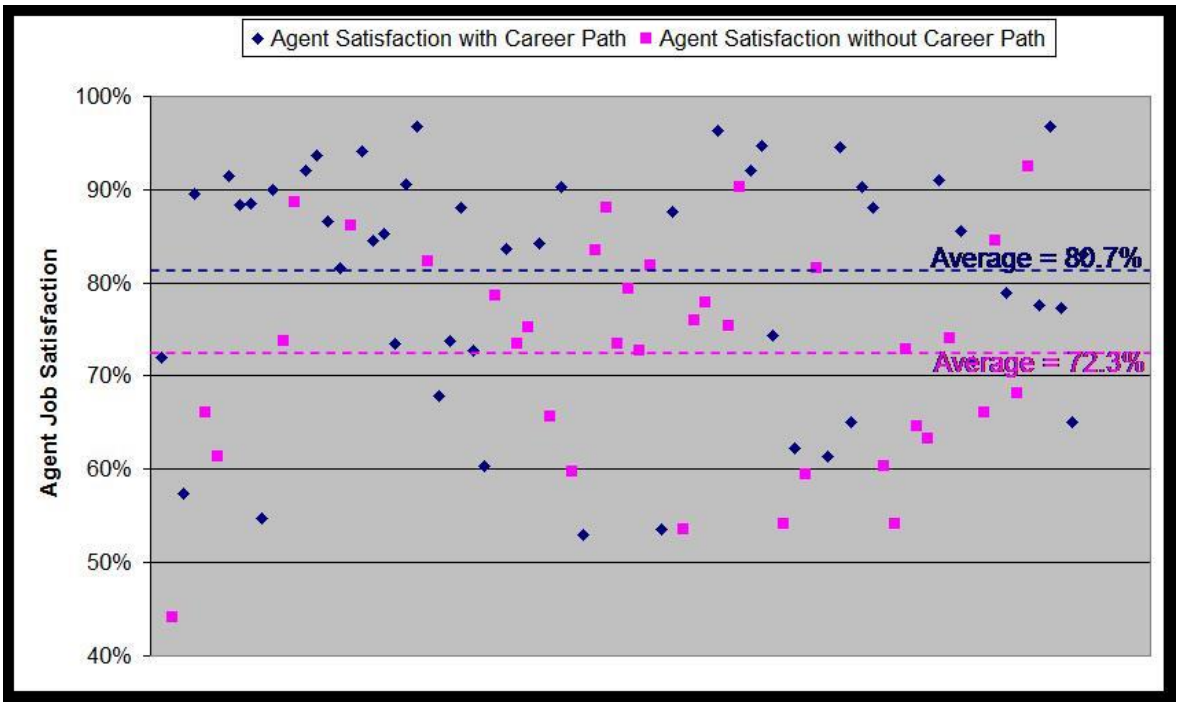

Benchmarking Ranges for Agent Job Satisfaction

MetricNet's benchmarking database shows that the average Agent Satisfaction for service desks worldwide is about 77%. That is, 77% of service desk agents are either satisfied or very satisfied with their jobs. This number varies from a low of 44%, to a high of 97%.

As mentioned above, Agent Job Satisfaction is usually measured through annual or semi-annual satisfaction surveys. These surveys typically contain five or fewer questions, and can be automated using simple online tools such as Survey Monkey. MetricNet recommends a five point scale, where a score of 5 represents Very Satisfied, and a score of 1 represents Very Dissatisfied. Given the impact of training, coaching, and career pathing on Agent Satisfaction, the survey should include a question about each of these. For example "How would you rank the quality of training you receive on the service desk?" The final, and most important question of the survey, should always be some variation of "How would you rate your overall job satisfaction on the service desk?"

Desktop Support Cost per Ticket

Cost per Ticket is the total monthly operating expense of Desktop Support divided by the monthly ticket volume. Operating expense includes the following components:

- Salaries and Benefits for Desktop Support Technicians
- Salaries and Benefits for Indirect Personnel (Team Leads, Supervisors, Workforce Schedulers, Dispatchers, QA/QC Personnel, Trainers, and Managers)
- Technology Expense (computers, software licensing fees, etc.)
- Telecom Expense
- Facilities Expense (office space, utilities, insurance, etc.)
- Travel, Training, and Office Supplies

As you might expect, the vast majority of costs for desktop support are personnel related. Figure 1 below shows the average breakdown of costs for North American desktop support organizations in 2010.

Figure 1: The Cost of Desktop Support

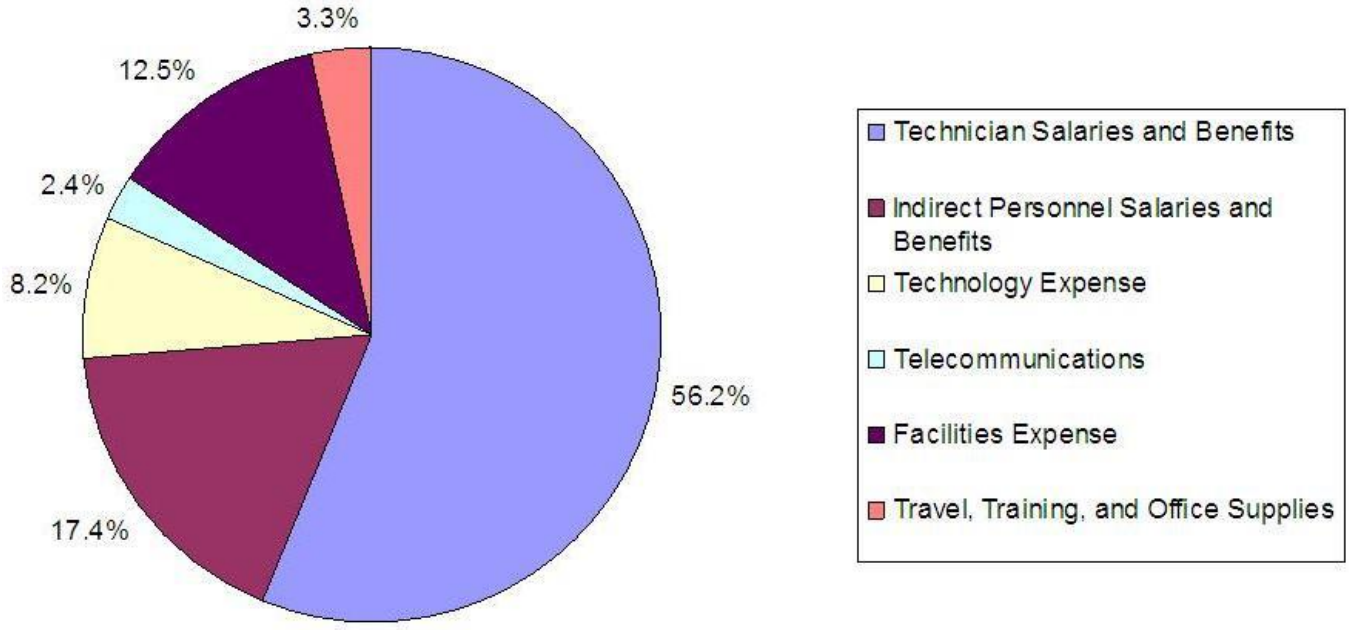

Why it's Important

Cost per Ticket, along with Customer Satisfaction, are often referred to as the *foundation metrics* in desktop support. They are the two most important metrics because ultimately everything boils

down to cost containment (as measured by Cost per Ticket), and quality of service (as measured by customer satisfaction).

In any service delivery organization, cost, or more accurately unit cost, is critically important. Cost per Ticket is a measure of how *efficiently* Desktop Support conducts its business. A higher than average Cost per Ticket is not necessarily a bad thing, particularly if accompanied by higher than average quality levels. Conversely, a low Cost per Ticket is not necessarily good, particularly if the low cost is achieved by sacrificing quality of service. Every Desktop Support organization should track and trend Cost per Ticket on a monthly basis.

Deconstructing Cost per Ticket

Desktop Support Tickets are comprised of incidents and service requests, as shown in Figure 2 below. Just as Cost per Contact gives us the unit cost for the level 1 service desk, Cost per Ticket, Cost per Incident and Cost per Service Request give us the unit costs for desktop support.

Figure 2: Tickets, Incidents and Service Requests

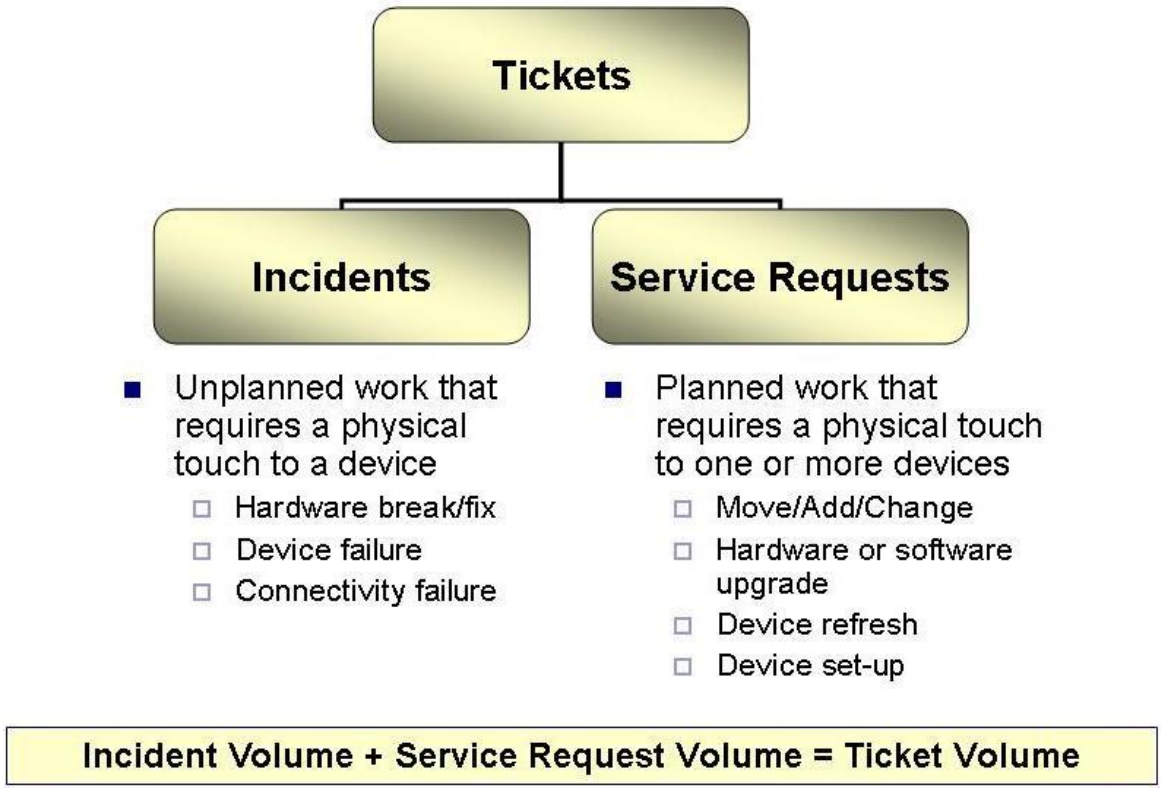

In general, incidents require less time to resolve than service requests, so the Cost per Incident is almost always less than the Cost per Service Request. Since the Cost per Ticket is a weighted average of the Cost per Incident and the Cost per Service Request, the mix of incidents and service requests will strongly influence the Cost per Ticket.

Let's assume, for example, that at Company ABC the Cost per Incident is $50, while the Cost per Service Request is $100. Additionally, 75% of ABC's tickets are incidents, while the remaining 25% of tickets are service requests. The Cost per Ticket can be calculated based upon a weighted average as follows: Cost per Ticket = ($50 X 75%) + ($100 X 25%) = $62.50.

Now let's take the example of Company XYZ who has the same Cost per Incident and Cost per Service Request as Company ABC, but whose mix of incidents and service requests is very different. Specifically, at XYZ only 40% of tickets are incidents, while the remaining 60% of tickets are service requests. The weighted average Cost per Ticket for XYZ can be calculated as follows: Cost per Ticket = ($50 X 40%) + ($100 X 60%) = $80.00.

So although ABC and XYZ have the exact same Cost per Incident and Cost per Service Request, their unique mix of incidents and service requests yields a very different Cost per Ticket. If both organizations were to handle 5,000 tickets per month, Company XYZ would spend $87,500 more per month on desktop support than Company ABC, simply because a larger percentage of their tickets are service requests!

Another key factor that influences the cost per ticket is the user population density. Desktop support technicians working in a high density user environment (e.g., a high-rise office building with lots of cubicles) are able to handle a larger volume of tickets per month than a technician supporting a work environment spread across a vast geographical area (think desktop support for a retail bank with hundreds of branches across the country).

Benchmarking Ranges for Cost per Ticket

The North American averages and ranges for Cost per Ticket, Cost per Incident, and Cost per Service Request in 2010 are shown below, in Figure 3.

Figure 3: North American Desktop Support Costs

Metric Type	Desktop Support KPI's	North American Statistics		
		Average	Min	Max
Cost	Cost per Ticket	$62	$27	$490
	Cost per Incident	$48	$19	$312
	Cost per Service Request	$113	$41	$556

The ranges on these metrics (min to max) are enormous, varying by more than an order of magnitude (10X) from min to max. The reasons for this were outlined above, and include the user population density, and the unique mix of incidents and service requests.

Agent Utilization

One goal of every business is to achieve the highest possible quality at the lowest possible cost. It stands to reason, therefore, that cost and quality should be measured on an ongoing basis. In fact, many would argue that cost and quality are the *only* two things that really matter in a service desk. In past articles MetricNet has discussed the importance of using metrics as a diagnostic tool to improve performance. So we have to ask ourselves, if cost per contact is one of the foundation metrics for the service desk, how can we affect it? How can we improve it? What are the primary levers we have to manage cost?

The service desk is a labor-intensive function. Agent salaries and benefits make up more than half of all costs for the average service desk. And when you consider the salaries and benefits for non-agents – e.g. supervisors, team leads, QA/QC, trainers, and workforce schedulers – approximately two thirds of all service desk costs are personnel related, as shown in Figure 1 below. Labor productivity then is the biggest lever we have to manage and control service desk costs.

Figure 1: Expense Breakdown for Global Service Desks

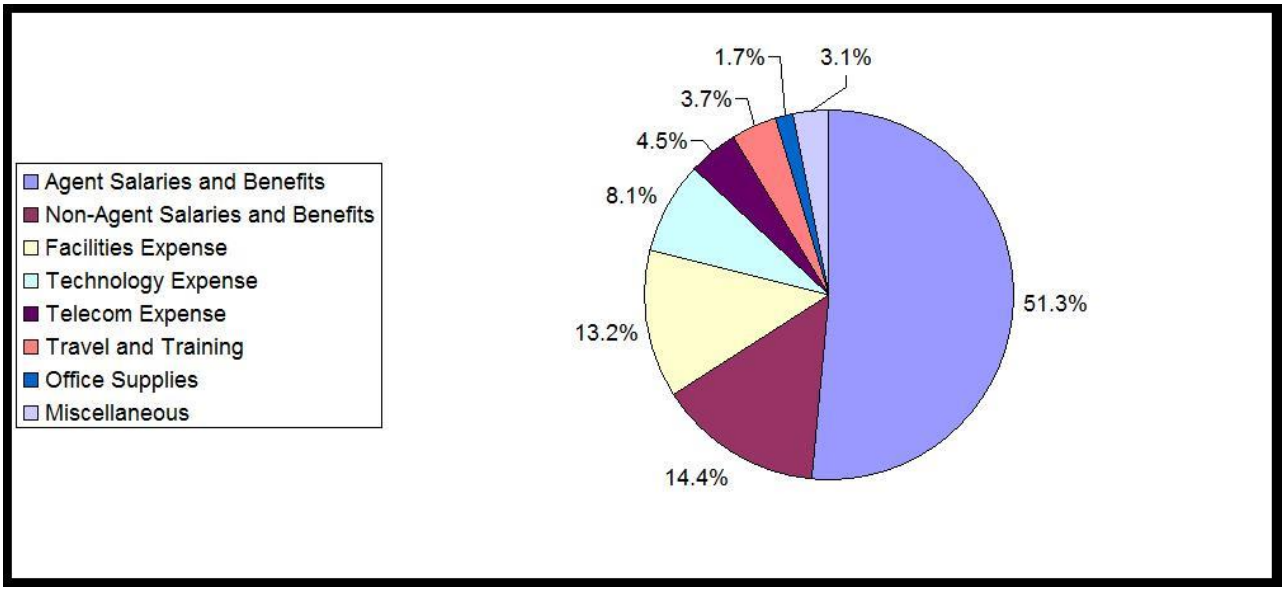

The best measure of labor productivity is agent utilization. Because agent salaries and benefits represent more than half of all service desk costs, if agent utilization is high, the cost per contact will be correspondingly low. Conversely, when agent utilization is low, agent costs, and hence cost per contact, will be correspondingly high. This is illustrated in Figure 2 below.

Figure 2: Agent Utilization vs. Cost per Contact

Just as world-class service desks are obsessive about maintaining high customer satisfaction levels, they are equally committed to keeping their costs in check. They do this primarily by maintaining tight control over agent utilization. This has the effect of minimizing cost per contact as illustrated above. That said, extremely high agent utilization rates can actually increase your costs by driving agent turnover rates higher. Whenever agent utilization rates approach 60% – 70 %, a service desk will experience relatively high agent turnover because they are pushing the agents too hard. Extremely high utilization leads to burnout, high turnover, and low morale in the service desk.

The formula for determining agent utilization is somewhat complicated. It factors in the hours in a work day, break times, vacation and sick days, training time and a number of other factors. But there is an easy way to approximate agent utilization without going to so much trouble:

$$\text{Agent Utilization} = \frac{(\text{Average number of calls handled by an agent in a month}) \text{ X } (\text{Average call handle time in minutes})}{(\text{Average number of days worked in a month}) \text{ X } (\text{Number of work hours in a day}) \text{ X } (60 \text{ minutes/hr})}$$

Let's say, for example that the agents in a particular service desk handle an average of 500 contacts per month at an average handle time of 10 minutes per contact. Additionally, these agents work an average of 21 days per month, and their work day is 7.5 hours after subtracting lunch and break times. The simplified utilization formula above would work out to the following:

$$\text{Agent Utilization} = \frac{(500 \text{ contacts/month}) \text{ X } (10 \text{ minutes/call})}{(21 \text{ working days per month}) \text{ X } (7.5 \text{ work hours per day}) \text{ X } (60 \text{ minutes/hr})} = 52.9\%$$

Once again, this is not a perfect measure of agent utilization, but it is quick and easy, and gets you within 5% of the true agent utilization figure.

Benchmark Ranges for Agent Utilization

MetricNet's benchmarking database shows that the average agent utilization for service desks worldwide is about 48%. This number varies widely, from a low of 22% to a high of 76%. Those at the low end of this range tend to be smaller service desks that are unable to achieve economies of scale, or service desks that are required to staff a back shift where few calls come in. Those at the high end of this range tend to be outsourcers who typically have good scale economies, and have a profit incentive to keep agent utilization rates on the high side.

Figure 3 below shows the distribution of agent utilization statistics for a representative cross-section of service desks worldwide.

Figure 3: Benchmarking Statistics for Agent Utilization

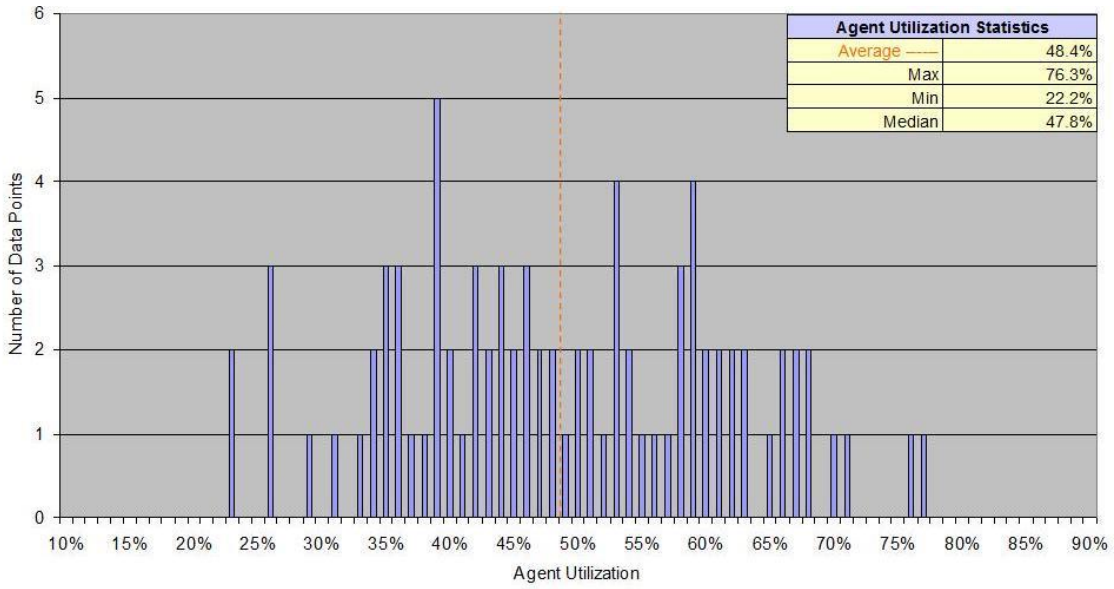

Tickets per User per Month

Tickets per User per Month is both a Service Desk and a Desktop Support metric. There are, however, important differences in the key drivers of these two metrics. So in this month's article we will focus on *Desktop Support* Tickets per User per Month.

As the name suggests, Tickets per User per Month is simply the total number of monthly tickets logged by desktop support divided by the number of users supported by desktop support.

For purposes of this discussion it is important to remember that desktop support tickets are comprised of both incidents and service requests. An Incident is typically unplanned work that requires the assistance of an on-site technician to resolve. Common examples include a desktop or laptop computer break/fix, a printer or server failure, connectivity problems, or any other issue that cannot be resolved remotely by the Level 1 Service Desk. By contrast, most Service Requests represent planned work. Among the most common Service Requests are Move's/Add's/Change's, hardware refresh/replacement, and device upgrades. Tickets represent the sum of all Incidents and Service Requests, as illustrated in Figure 1 below.

Figure 1: Tickets, Incidents and Service Requests

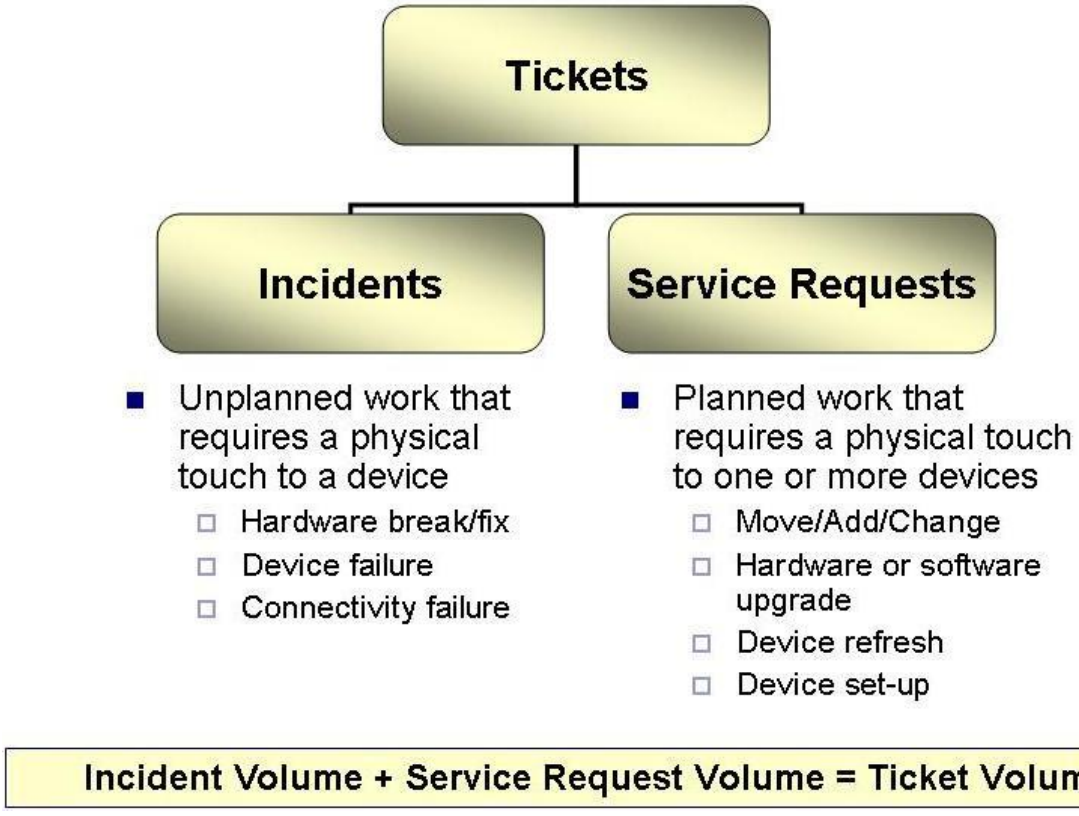

Why it's Important

Tickets are the primary unit of work in desktop support. As such, ticket volume will drive the headcount of technicians needed by an organization. A common misperception in desktop support is that the user population alone will define the number of technicians needed. This approach wrongly assumes that the ratio of desktop support technicians to the number of users is fixed. For example, 12.5 desktop support technicians are needed for every 1,000 users. The error in this approach is that no two user populations have the same needs, and therefore no two user populations generate the same workload. As such, staffing decisions in desktop support should be based upon *workload*, not user population. With this in mind, it is easy to see why two organizations with exactly the same headcount may require very different staffing levels in desktop support.

Benchmark Data for Tickets per User per Month

The number of tickets generated by an organization is driven by numerous factors including the average age of devices supported, the mix of laptop and desktop computers, the number of remote users, the number of mobile devices, the refresh rate of devices, the standardization (or lack thereof) of the IT environment, and the degree of virtualization.

Figures 2, 3, and 4 below show just how dramatically the incident and service request volume can vary from company to company, and from industry to industry. Average monthly ticket volumes range from a low of just 0.41 tickets per seat per month in health care, to a high of 0.99 tickets per seat per month in Financial Services.

Figure 2: Monthly Ticket Volume by Industry

Key Performance Indicator	Statistic	Financial Services	High Tech	Equipment Manufacturing	Telecom	Business Services	Health Care	Energy Utilities
Incidents per Seat per Month	Average	0.67	0.56	0.39	0.43	0.72	0.30	0.36
	Min	0.19	0.14	0.11	0.12	0.22	0.12	0.07
	Max	1.95	1.82	1.24	1.56	2.07	0.65	1.01
Service Requests per Seat per Month	Average	0.32	0.42	0.22	0.31	0.42	0.11	0.20
	Min	0.10	0.29	0.14	0.13	0.14	0.05	0.10
	Max	1.20	1.41	0.62	0.94	1.44	0.36	0.59
Total Tickets per Seat per Month	Average	0.99	0.98	0.61	0.74	1.14	0.41	0.56
	Min	0.29	0.43	0.25	0.25	0.36	0.17	0.17
	Max	3.15	3.23	1.86	2.50	3.51	1.01	1.60

Figure 3: Monthly Incident Volume by Industry

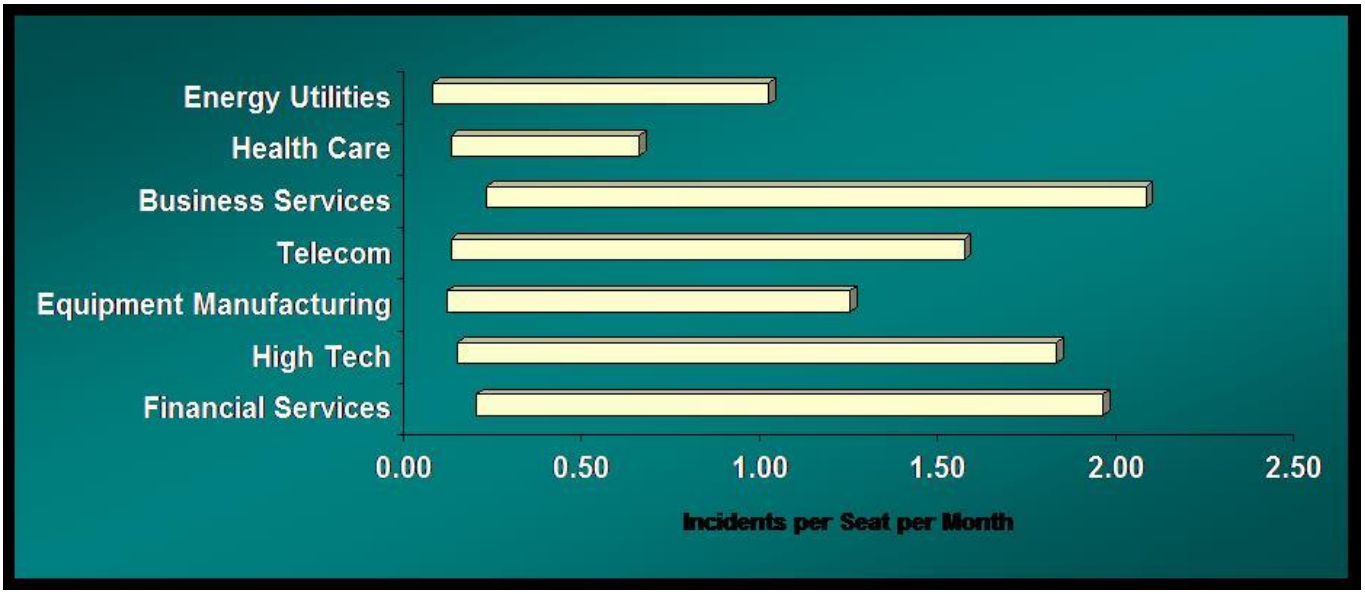

Figure 4: Monthly Service Request Volume by Industry

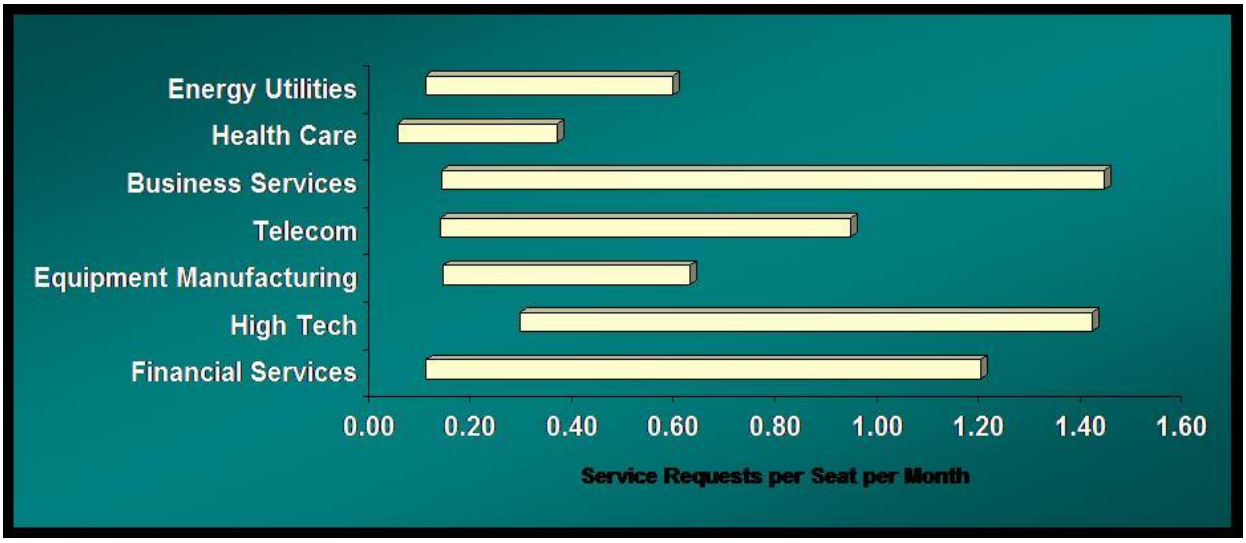

Since ticket volume drives technician headcount, one would also expect to see a wide variation in headcount requirements from company to company, and from industry to industry. Figures 5 and 6 below show that the average desktop technician headcount can range from a low of just 5.4 technicians per 1,000 seats in health care, to a high of 28.4 technicians per 1,000 seats in a high tech company.

Figure 5: Desktop Support Technician Headcount Ranges by Industry

Key Performance Indicator	Statistic	Financial Services	High Tech	Equipment Manufacturing	Telecom	Business Services	Health Care	Energy Utilities
Desktop Technicians per 1,000 Seats Supported	Average	21.9	28.4	12.7	15.5	27.0	5.4	7.7
	Min	3.4	7.5	2.4	2.6	3.6	1.1	1.2
	Max	67.3	90.1	40.5	40.4	55.2	12.3	16.3

Figure 6: Desktop Support Technician Headcount Ranges by Industry

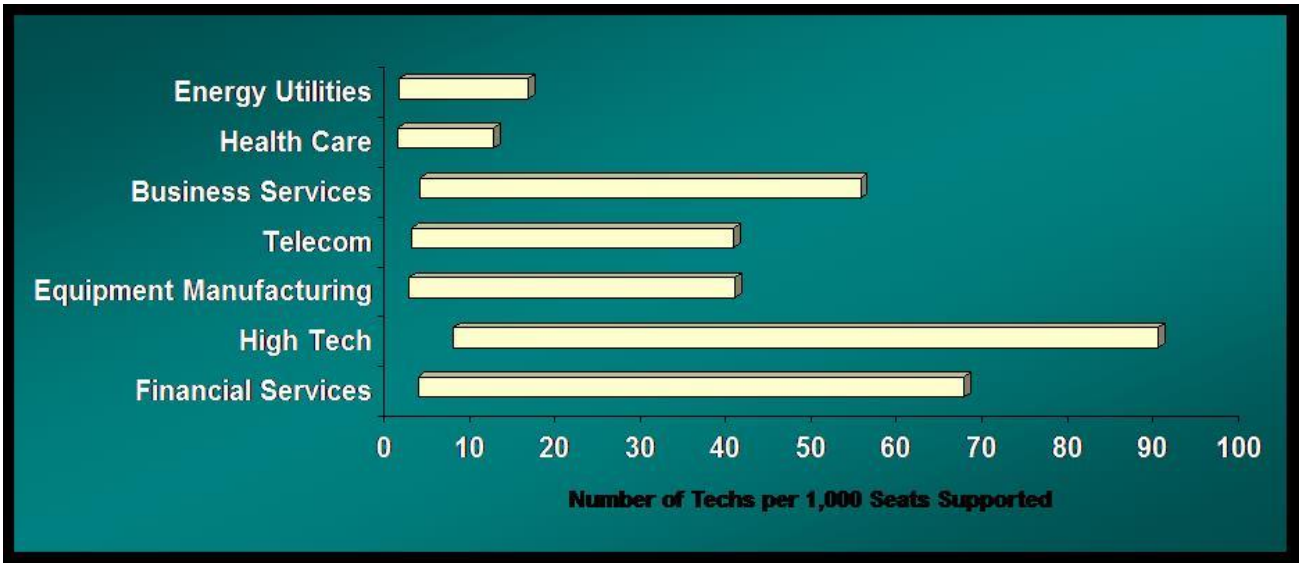

Call Abandonment Rate

An abandoned call is one where the caller hangs up before being connected to a live agent in the service desk. Call abandonment rate is the number of abandoned calls divided by all calls offered to the service desk, and is one of the most widely tracked metrics in the service desk industry. Virtually every service desk with an ACD has the ability to track this metric.

With the advent of chat, and the increasing number of support requests handled through chat, abandonment rate has come to signify both abandoned calls and abandoned chat sessions – i.e., chat sessions that are initiated by the user but abandoned before an agent responds. Ideally chat and voice abandonment rates should be tracked separately, as any significant imbalance between the two abandonment rates could indicate an imbalance in staffing between the two channels. For purposes of this article, abandonment rate will refer to abandoned voice calls, but the principles discussed apply equally to abandoned calls and chats.

Why it's Important

Many service level agreements include an abandonment rate target. Although a low abandonment rate is a worthy objective, many service desks go too far in trying to reduce abandoned calls. Those who pursue this strategy typically believe that a low abandonment rate is a prerequisite for achieving high customer satisfaction levels. The truth, however, is that a low abandonment rate will not necessarily lead to higher levels of customer satisfaction. As discussed in a prior *Metric of the Month*, the real driver of customer satisfaction is first contact resolution rate. So while customers may be willing to forgive an occasional abandoned call, they are far less forgiving of calls that are not resolved on first contact.

Figure 1 below shows the relationship between abandoned calls and customer satisfaction for a representative cross section of North American service desks. At low abandonment rates the relationship between customer satisfaction and abandonment rate is almost flat, meaning that abandonment rate has very little impact on customer satisfaction levels. It is only when the abandonment rate increases above seven or eight percent that customer satisfaction begins to drop off significantly.

Figure 1: Call Abandonment Rate vs. Customer Satisfaction

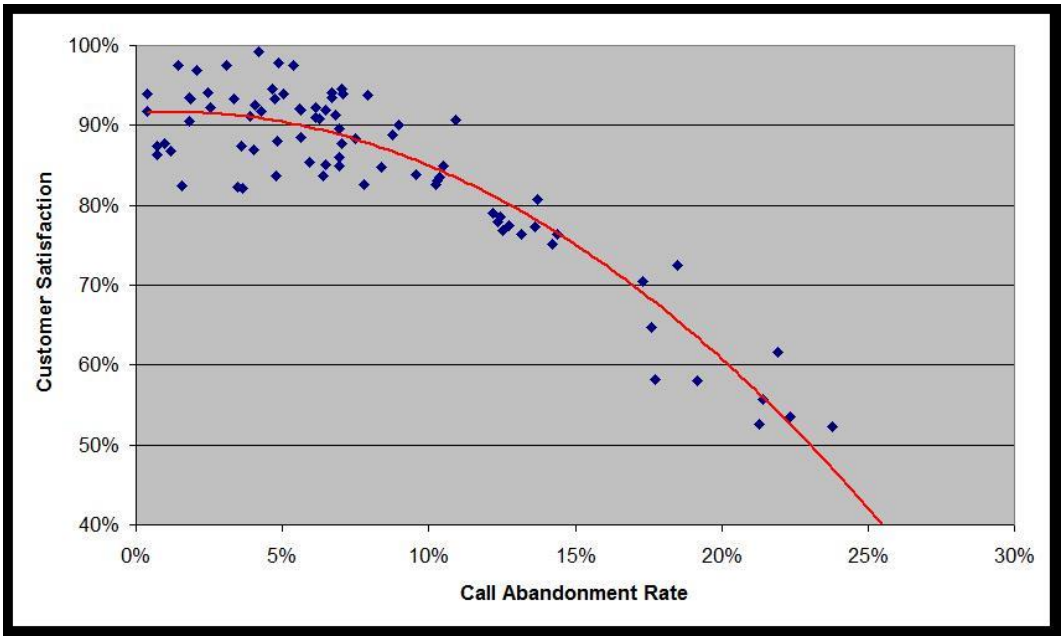

While low abandonment rates have virtually no impact on customer satisfaction, they *do* have a direct impact on cost per contact. In fact, the lower the abandonment rate is, the higher the cost per contact will be. Figure 2 below illustrates this phenomenon.

Figure 2: Call Abandonment Rate vs. Cost per Contact

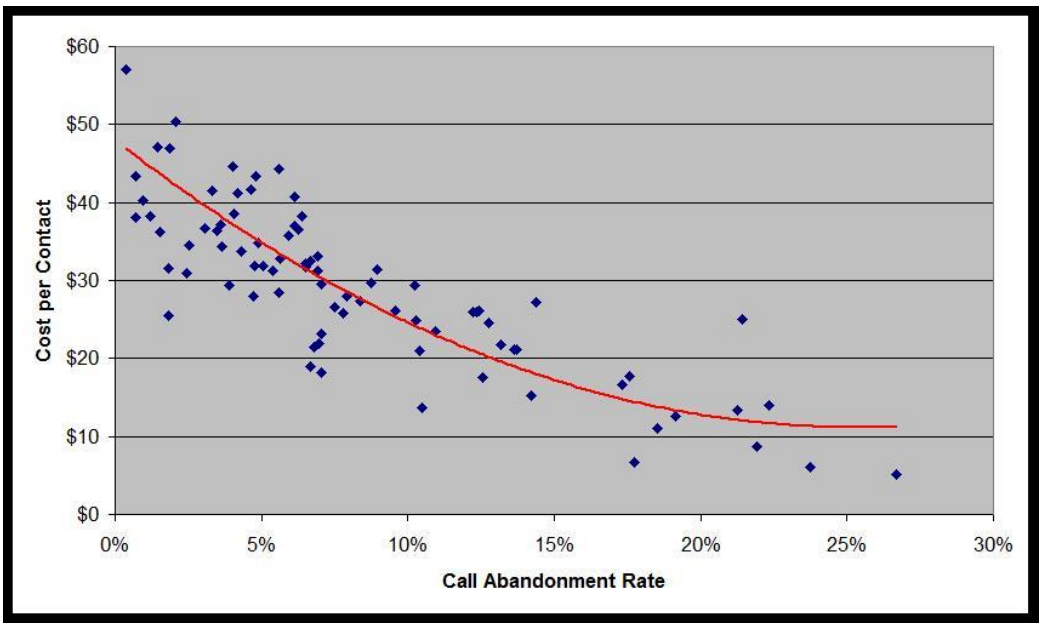

So why do costs go up as the abandonment rate goes down? Because more agent headcount is required to achieve lower abandonment rates, and as headcount increases so too does the cost of support. Let's say, for example, that a service desk with 10 full time agents has an average abandonment rate of 8%. To reduce the abandonment rate to 4% would require a full time agent headcount of approximately 13. That's a 30% increase in headcount, for a 4 percentage point

decrease in abandonment rate. Is that worth it? Apparently not, because the added headcount increases costs significantly, but produces very little benefit in terms of higher customer satisfaction.

This suggests that there is an optimal range for call abandonment rate. If you go much above this range customer satisfaction will drop off fairly quickly, and if you go much below the range costs will rapidly climb. Benchmarking data shows that the optimal range for abandoned calls is between 4% and 7%. Operating within this range will keep a service desk off the steepest part of the cost curve at low abandonment rates, and also prevent a dramatic drop off in customer satisfaction at higher abandonment rates.

Benchmark Data for Call Abandonment Rate

Industry data from MetricNet's benchmarking database shows that the average call abandonment rate for all service desks is 8.7%, which is slightly above the optimal range. Additionally, the number of service desks operating outside the optimal range is somewhat surprising. This suggests that service desks operating at low call abandonment rates may be unaware of the heavy price they pay for this performance. Conversely, service desks operating at high abandonment rates may be unaware of the adverse impact their abandonment rate has on customer satisfaction.

Figure 3: Industry Benchmarks for Call Abandonment Rate

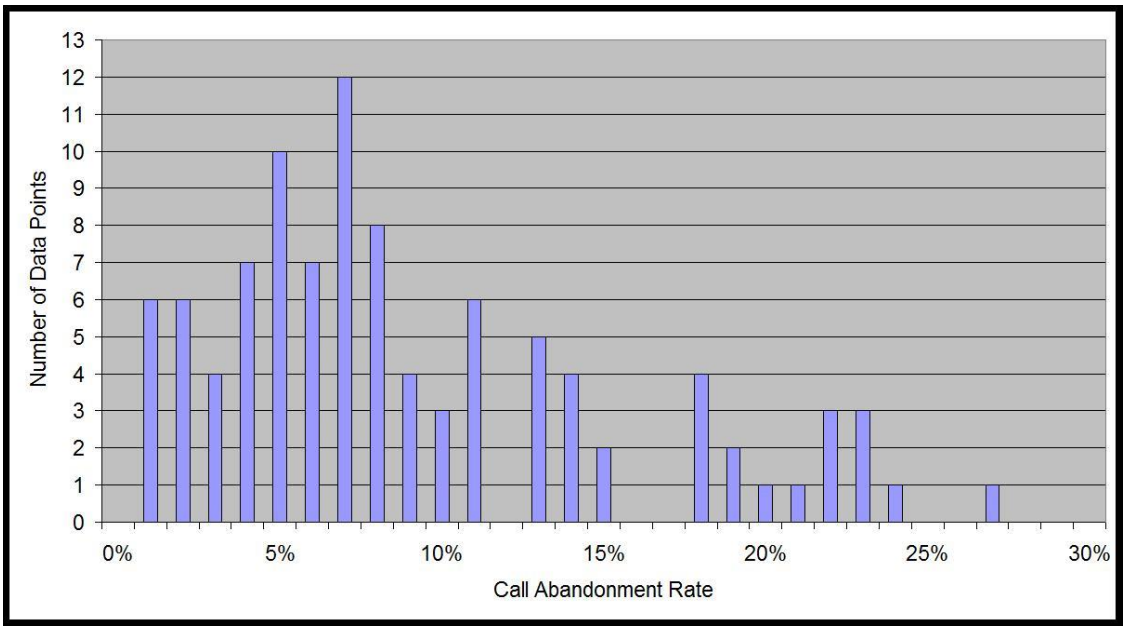

Contact Handle Time

Contact handle time is the average time that an agent spends on an inbound contact, including talk time, chat time, wrap time, and after call or after chat work time (ACW). For non-live contacts, such as email, voice mail, and faxes, the contact handle time is the average time that an Agent spends working on the contact before escalating or closing out the ticket.

Please note that contact handle time is not the same as mean time to resolve (MTTR). MTTR, sometimes called cycle time, is the average elapsed time from the beginning of an incident or service request, until the ticket associated with the incident or service request is closed out. So while the total handle time for an incident may be only 10 minutes, unless the ticket is resolved on first contact the MTTR will be longer than the handle time. In fact, the MTTR can be substantially longer than the handle time depending upon how much after call work is required to resolve the issue, and whether the ticket is escalated to another level of support for resolution.

Why it's Important

A contact is the basic unit of work in a service desk. Contact handle time, therefore, represents the amount of labor required to complete one unit of work. Additionally it is an indirect measure of contact complexity. As an example, a typical password reset has a handle time of one to three minutes. By contrast, a contact for a proprietary business application such as Oracle or SAP may have a handle time of 15 minutes or more.

Contact handle time has a direct impact on, and is directly impacted by several other service desk metrics. One productivity metric that is strongly impacted by handle time is the number of contacts that an agent handles in a month. For a 10 minute handle time, the average agent handles about 500 contacts per month. In a more complex environment, where the average handle time is 15 minutes, a typical agent might handle only 300 or so contacts per month. This, in turn, has headcount implications. A service desk that takes 5,000 inbound contacts per month with an average handle time of 15 minutes will require more agents than a service desk that takes 5,000 contacts per month with an average handle time of 10 minutes.

First contact resolution and first level resolution also have an impact on contact handle time. In general, handle times increase as the first contact and first level resolution rates increase. This stands to reason because it is the more difficult/complex contacts that are typically dispatched and/or escalated to higher levels of support, and an increase in FCR and FLR is an indication that the service desk is resolving a more complex set of contacts (read longer handle times) at level 1.

Since contact handle time is a proxy for complexity, service desks with longer handle times generally require more experienced agents. To continue with our earlier example, an agent that is expected to reset passwords and support basic desktop applications such as Microsoft Office does not need the level of experience or expertise that would be required of an agent that is expected to resolve incidents related to more complex business applications. This of course has

implications for training. Inasmuch as longer handle times represent more complex incidents, agents supporting these environments require more training and experience to be successful.

Perhaps the most significant implication of handle time is the impact it has on cost per contact. For obvious reasons, cost per contact increases as handle time increases. Furthermore, this is one of the few linear relationships in IT support. As such, projecting cost per contact vs. handle time can be done very accurately and is a relatively straightforward exercise.

Benchmark Data for Contact Handle Time

Industry data from MetricNet's benchmarking database shows that the average contact handle time for all service desks is just under 10 minutes.

Figure 1: Industry Benchmarks for Average Contact Handle Time

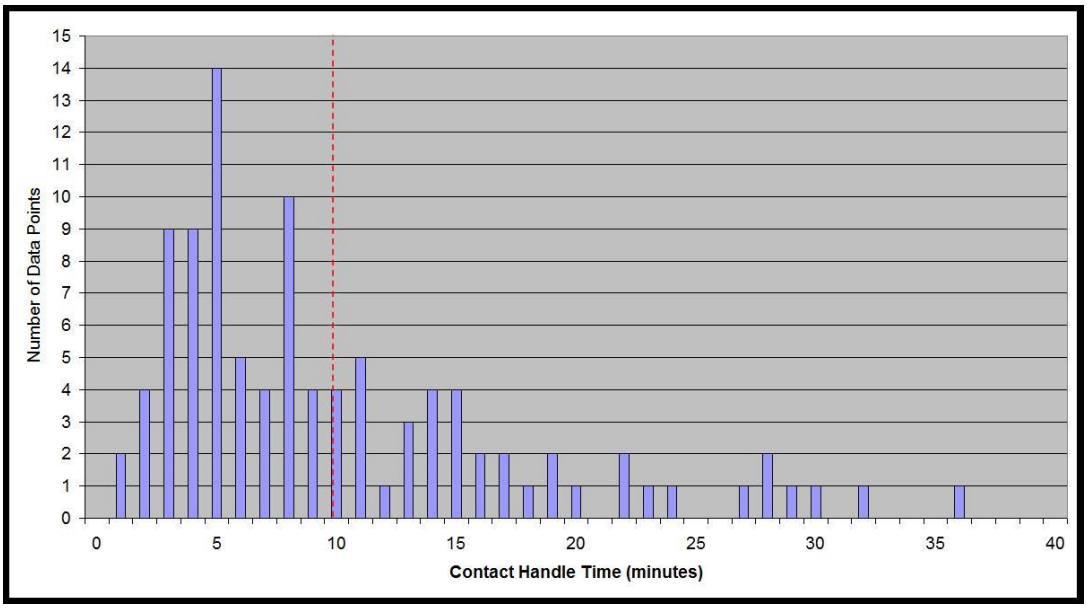

Figure 2 below shows the correlation between cost per contact and handle time. Although the relationship is linear, there is a fairly large variance around the mean (red line). This is due to differences in wage rates in different geographies, as well as the higher salaries paid to agents in the longer handle time (more complex) environments.

Figure 2: Cost per Contact vs. Contact Handle Time

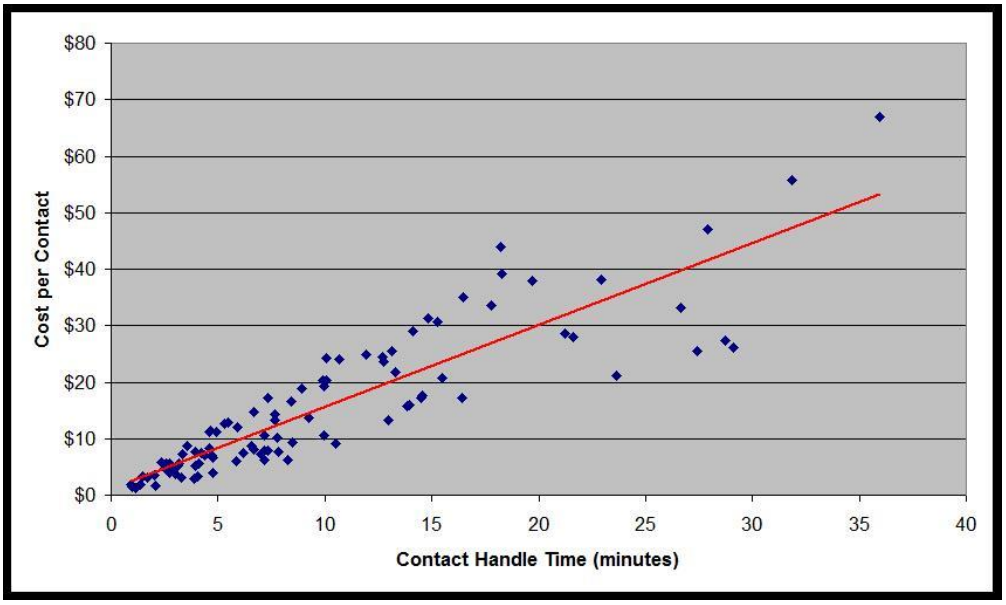
$80
$70
$60
$50
$40
$30
$20
$10
$0
Cost per Contact
0
5
10
15
20
25
30
35
40
Contact Handle Time (minutes)

Agent to Supervisor Ratio

The Agent to Supervisor Ratio is simply the number of front-line agents divided by the number of supervisors for a service desk. It is a measure of management span of control, and managerial efficiency. The comparable metric for desktop support is the number of technicians divided by the number of supervisors in desktop support. In this article we will use service desk examples and benchmarking data, but the principles discussed apply equally to the ratio of technicians to supervisors in desktop support.

Why it's Important

Service desk managers and supervisors frequently ask me about the proper ratio of agents to supervisors. Should it be 5 to 1? 10 to 1? 20 to 1? Like most KPIs there are tradeoffs involved with this metric. If the ratio is too high the management span of control is too broad, and agents can be working without the proper level of oversight and supervision. This, in turn, can lead to a multitude of issues ranging from low morale, to inadequate training, coaching, and feedback. By contrast, a ratio that is too low is an indication that a service desk is "top heavy"; it has too many supervisors for the number of agents. This, in turn, leads to higher costs, and specifically higher cost per contact.

There are at least three techniques that are commonly used to determine the proper ratio of agents to supervisors. The first technique is a bottom-up modeling approach that catalogs all of the duties and responsibilities of a supervisor, and then assigns a time value to each responsibility. A typical supervisor, for example, might be responsible for team meetings, agent coaching, handling overflow and/or problematic calls, team meetings, working with vendors, managing or participating in projects such as ITIL training, and numerous other duties. The time commitment required for these activities can be estimated and then summed up to determine the monthly supervisory workload for the desk.

Let's say, for example, that a particular service desk has 65 agents and an estimated supervisory workload of 1,200 hours per month. Since there are approximately 172 working hours in a month, we can estimate the supervisory headcount to be about seven (1,200 supervisory hours per month ÷ 172 work hours per month = 7 supervisors). The agent to supervisor ratio for this desk would therefore be 9.3 (65 agents ÷ 7 supervisors = 9.3 Agents per Supervisor).

The second approach relies upon industry benchmarks to provide guidance on the proper ratio of agents to supervisors, while the third and final approach for determining this metric is to look at confirming metrics such as agent job satisfaction. Both of these approaches are discussed below.

Benchmark Data for Agent to Supervisor Ratio

Industry data from MetricNet's benchmarking database shows that the average ratio of agents to supervisors is 8.5, but ranges widely, from a high of 21 agents per supervisor to a low of just 1.8 agents per supervisor.

Figure 1: Industry Benchmarks for Agent to Supervisor Ratio

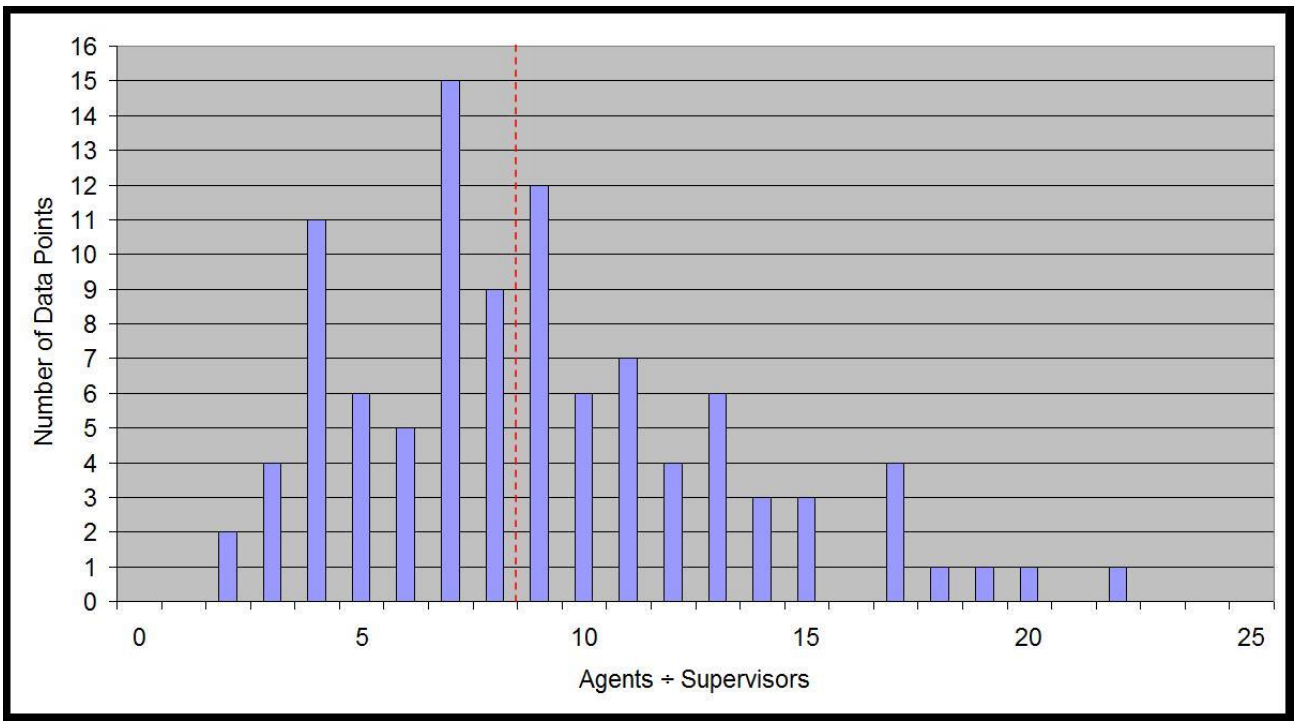

Figure 2 below shows the relationship between agent job satisfaction and the agent to supervisor ratio. As the ratio decreases, agent job satisfaction increases, presumably because adequate supervision contributes to the agent's morale and well-being. However, diminishing returns come into play at lower ratios, and there appears to be very little benefit in reducing the ratio to less than five agents for every supervisor. Finally, it is worth pointing out once again that the tradeoff here is one of cost vs. the benefits of a lower ratio. While high levels of agent morale and job satisfaction are just one of the benefits of a low agent to supervisor ratio, the cost per contact will continue to increase as the ratio decreases.

Figure 2: Agent to Supervisor Ratio vs. Agent Job Satisfaction

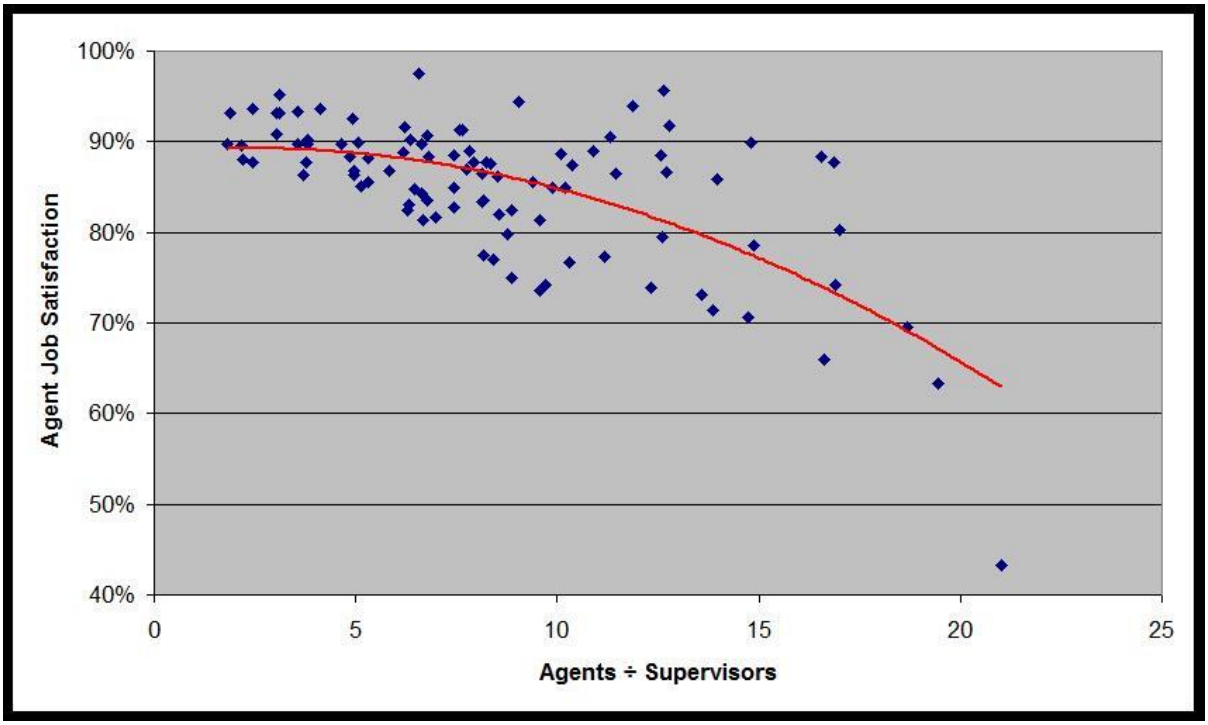
100%
90%
80%
70%
60%
50%
40%
Agent Job Satisfaction
0
5
10
15
20
25
Agents ÷ Supervisors

Customer Satisfaction

Customer Satisfaction is the percentage of customers that are either satisfied or very satisfied with the quality of support they receive. It is equally applicable to the service desk and desktop support, and is the single most important measure of quality for a support organization. In this article we will use service desk examples and benchmarking data, but the principles discussed apply equally to desktop support.

Cost vs. Quality

Any discussion of customer satisfaction must take place in the context of cost. Customer Satisfaction and cost per contact are the yin and yang of a service organization. These metrics are in constant tension, and every service organization grapples with how to strike an appropriate balance between the two. With no restraints on spending, it is relatively easy for a service desk to "spend its way" to high customer satisfaction. Conversely, if customer satisfaction is not an issue, a service desk can reduce its costs almost indefinitely.

Every support organization does, of course, operate with both spending limitations and quality standards. The objective, therefore, is to achieve the highest possible quality within the budgetary constraints placed on an organization. Some service desks that operate under severe budgetary restrictions are doing quite well when they achieve a 75% customer satisfaction score. By contrast, a service desk with a relatively generous budget that achieves an 85% customer satisfaction score may not be doing so well, despite having a higher absolute customer satisfaction score. The point is that cost matters, and customer satisfaction must be interpreted in light of the spending constraints that a service desk operates under.

Figure 1 below shows the relationship between customer satisfaction and cost per contact for a representative cross-section of global service desks. The tradeoff between cost and quality is apparent, but at any given spending level the more effective support organizations are able to achieve higher levels of customer satisfaction.

Figure 1: Cost per Contact vs. Customer Satisfaction

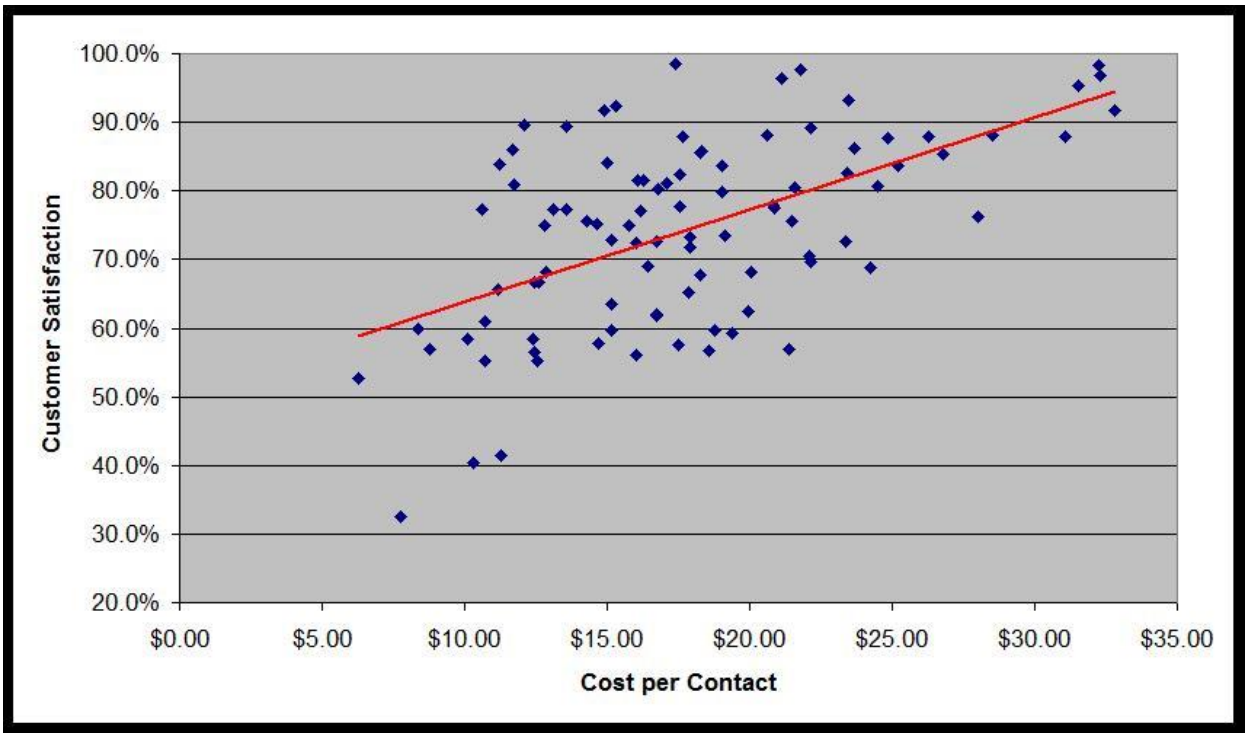

Cost per contact and customer satisfaction are often referred to as the foundation metrics in support because everything boils down to these two metrics. This insight is both enlightening and liberating. Viewed through the twin lenses of "will it reduce cost" and "will it improve customer satisfaction", management decisions are often greatly simplified by answering these two questions.

How It's Measured

There are almost as many different ways to measure customer satisfaction as there are service desks that track the metric. I have seen surveys that contain as few as one question, and surveys that contain as many as 40 questions. I have seen multiple choice, fill-in-the-blank, and interview style surveys. I have seen scoring systems that offer as few as two choices per question, and as many as twelve choices per question. The result is that customer satisfaction has the greatest variability of any metric in the service desk.

When you think about the number of different survey questions, survey types, scoring scales, and survey populations, it is easy to understand why survey bias and invalid survey methodologies are so common in the industry. HDI's Customer Satisfaction Index has addressed this shortcoming, and presents a valid, standardized approach to measuring customer satisfaction for all support organizations.

Benchmark Data for Customer Satisfaction

Industry data from MetricNet's global benchmarking database shows that the average customer satisfaction is 74.4%, but ranges widely, from a high of 98.6% to a low of 32.6%.

Figure 2: Industry Benchmarks for Customer Satisfaction

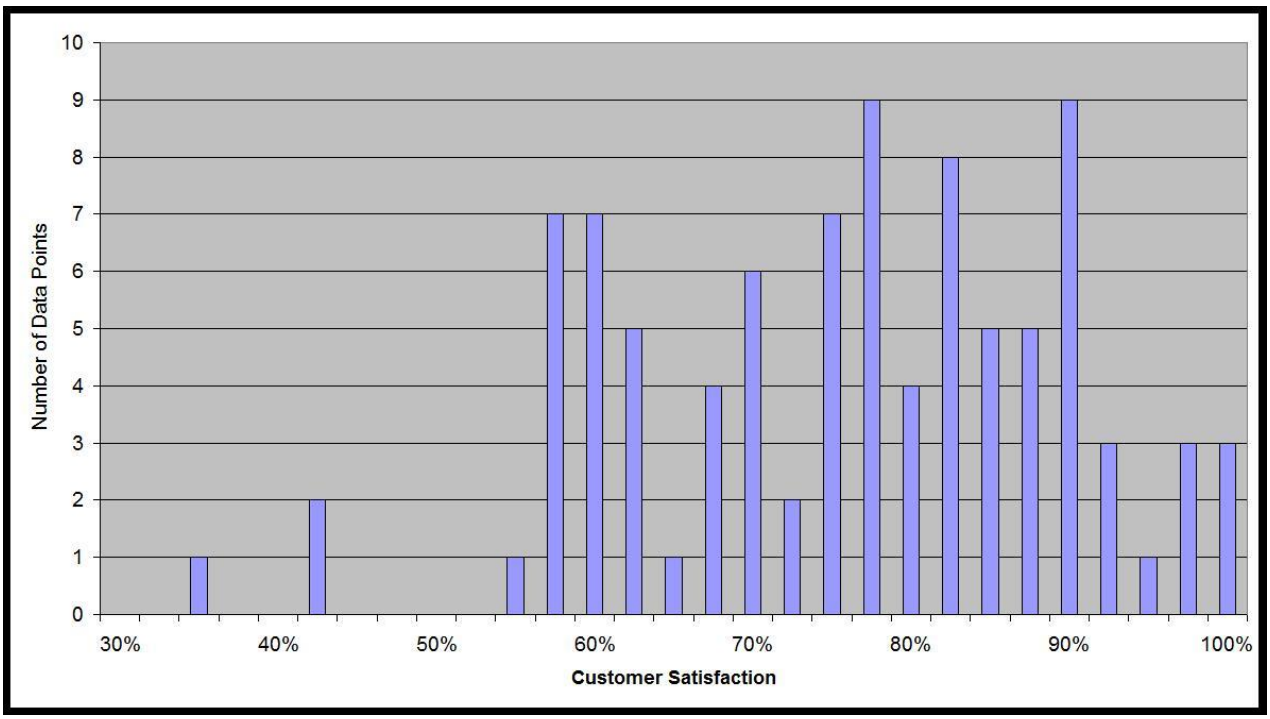

Finally, it is worth pointing out as we did in the August 2011 Metric of the Month that the biggest driver of customer satisfaction is first contact resolution, as shown in Figure 3 below. The implication is that the surest way to improve customer satisfaction is to increase the first contact resolution rate.

Figure 3: First Contact Resolution vs. Customer Satisfaction

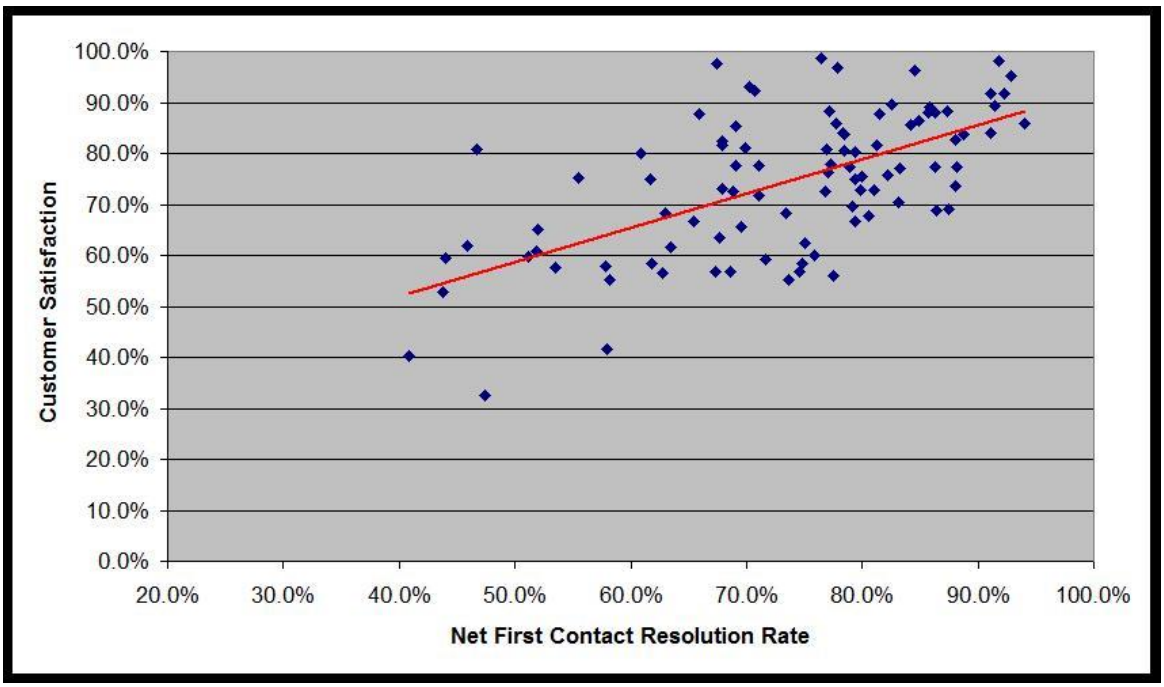

Mean Time to Resolve

Mean Time to Resolve (MTTR) is a service level metric for desktop support that measures the average elapsed time from when an incident is reported until the incident is resolved. It is typically measured in hours, and refers to *business hours*, not clock hours. A desktop incident that is reported at 4:00 pm on a Friday and closed out at 4:00 pm the following Monday, for example, will have a resolution time of eight business hours, *not* 72 clock hours. Most incident management systems can easily track MTTR.

Please note that we make a distinction between incidents and service requests. A desktop incident is typically unplanned work that requires the assistance of an on-site technician to resolve. Common examples include break/fix requests for a laptop computer, a printer or server failure, connectivity problems, or other issues that cannot be resolved remotely by the level 1 service desk. By contrast, most desktop service requests represent planned work. Among the most common desktop service requests are move/add/change's, hardware refresh/replacement, and device upgrades. MTTR as discussed in this article refers specifically to incidents, not service requests.

Why It's Important

As you know from prior Metric of the Month articles, service levels at level 1, including average speed of answer and call abandonment rate, are relatively unimportant. They have little if any influence on customer satisfaction. The same, however, cannot be said of service levels for desktop support. In fact, MTTR is one of the key drivers of customer satisfaction for desktop support. This makes sense, as a user may be completely down, or forced to use workarounds until their incident has been resolved. This, in turn, has a significant impact on their overall satisfaction with desktop support.

Figure 1 below shows the relationship between customer satisfaction and incident MTTR for a representative cross-section of global desktop support groups. The strong correlation between MTTR and customer satisfaction is readily apparent.

Figure 1: Incident MTTR vs. Customer Satisfaction

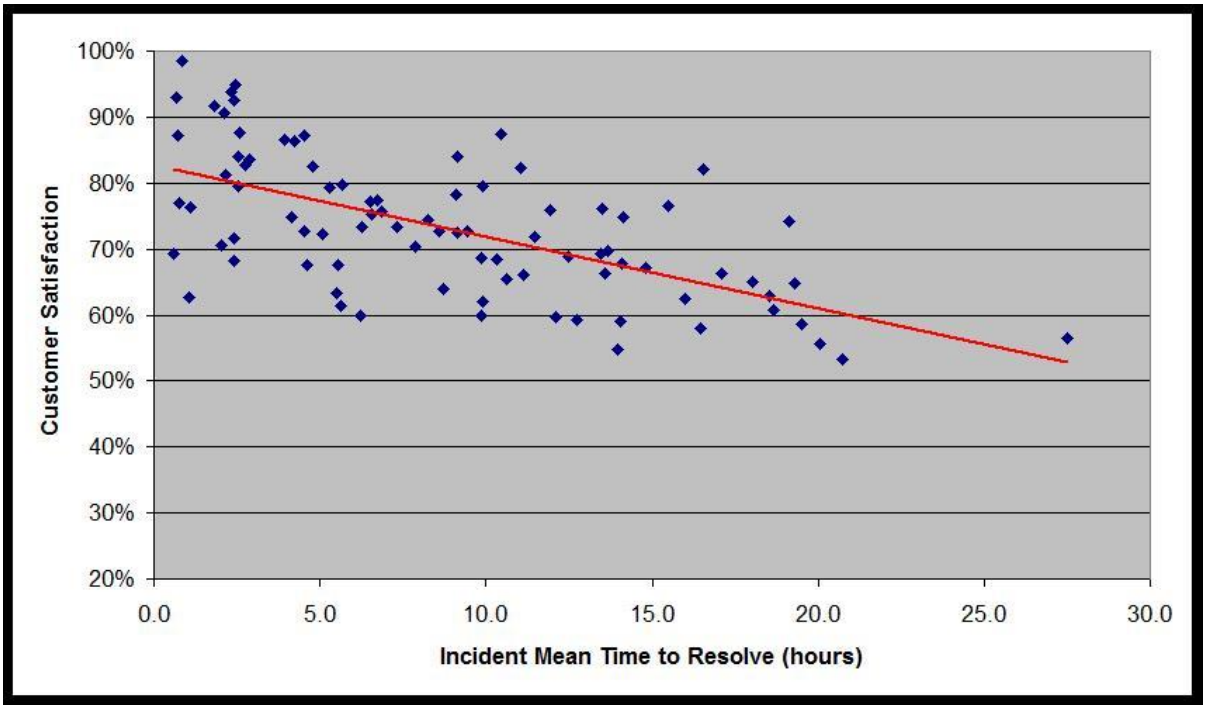

Benchmark Data for Incident MTTR

Industry data from MetricNet's global benchmarking database shows that the average incident MTTR is 8.85 business hours, but ranges widely, from a high of 27.5 hours to a low of 0.6 hours (Figure 2 below). This huge variation is driven largely by the user population density. In a high density user environment (e.g., a high-rise office building with lots of cubicles), the technician travel time to and from the site of an incident is generally short, e.g. less than 10 minutes. This results in shorter MTTR's. By contrast, technician travel time for users who are spread out over a broader geographical area (think desktop support for a retail bank with many branches in a given region) is often significantly greater, and will increase the MTTR accordingly.

Figure 2: Industry Benchmarks for Mean Time to Resolve

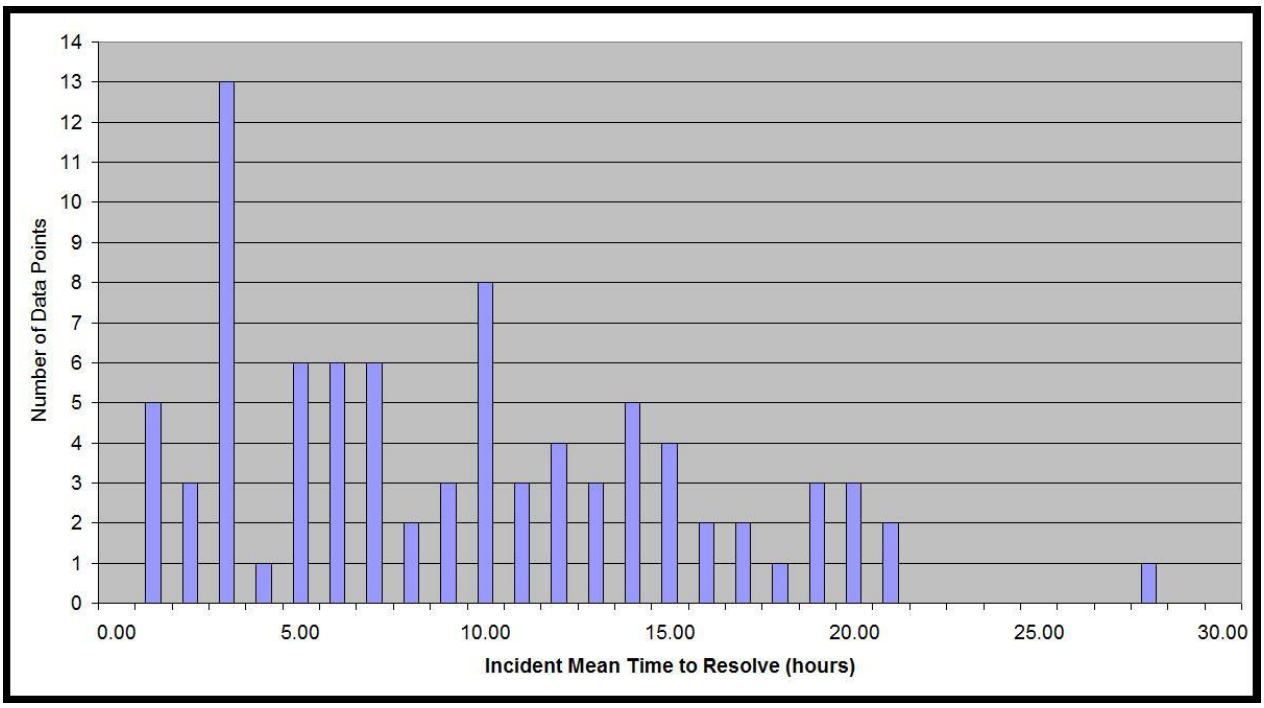

Inasmuch as customer satisfaction is driven by MTTR, many desktop support organizations take steps to manage and minimize this metric. Although the user population density, and hence the travel time per incident cannot be controlled, other factors affecting MTTR *can* be managed. These include maximizing the first visit resolution rate (comparable to first contact resolution rate at level 1), and routing desktop technicians in real time. This latter technique allows an organization to manage the incident queue by dispatching technicians based on the proximity and geographic clustering of incidents rather than on a first-in-first-out (FIFO) basis, as is common in the industry. This has been shown to significantly reduce MTTR for desktop incidents.

Agent Training Hours

There are two commonly tracked service desk metrics for agent training: new agent training hours and annual agent training hours. New agent training measures the number of training hours, including classroom, computer-based-training, self-study, direct coaching, and on-the-job training, that a new agent receives before they are allowed to handle customer contacts independently. Annual agent training, by contrast, includes any ongoing training that an agent receives after their new agent training has been completed. Annual agent training does not include routine team meetings, shift handoffs, or other activities that do not involve formal training.

Most support organizations will tell you that they don't do as much training as they would like. The most common reasons cited for the shortfall are budget and time constraints. But given the overwhelmingly positive impact that training can have on an organization, it's no surprise that the staunchest defenders of the training budget tend to have the best performing support organizations.

Why It's Important

Agent training is one of the few metrics that a support organization has complete control over. Moreover, it is a preventive/proactive metric because higher training hours are positively correlated with several other important metrics on the service desk. For example, agent training is one of the key drivers of first contact resolution rate (Figure 1). And as we know from prior Metric of the Month articles, first contract resolution is the single biggest driver of customer satisfaction. This is why support organizations looking to improve their customer satisfaction scores will oftentimes increase their agent training hours in a bid to increase first contact resolution rates, which then drives higher customer satisfaction levels.

Figure 1: Agent Training Hours vs. First Contact Resolution Rate

Training hours also have a direct impact on agent job satisfaction, and hence the overall morale of an organization (Figure 2). As many experienced service and support professionals can attest, maintaining good morale in the workforce can be a challenge. Clearly one of the benefits of providing agents with an abundance of training opportunities is that it can help to boost and sustain high levels of agent job satisfaction. This, in turn, has the effect of reducing agent turnover and absenteeism.

Figure 2: Agent Training Hours vs. Agent Job Satisfaction

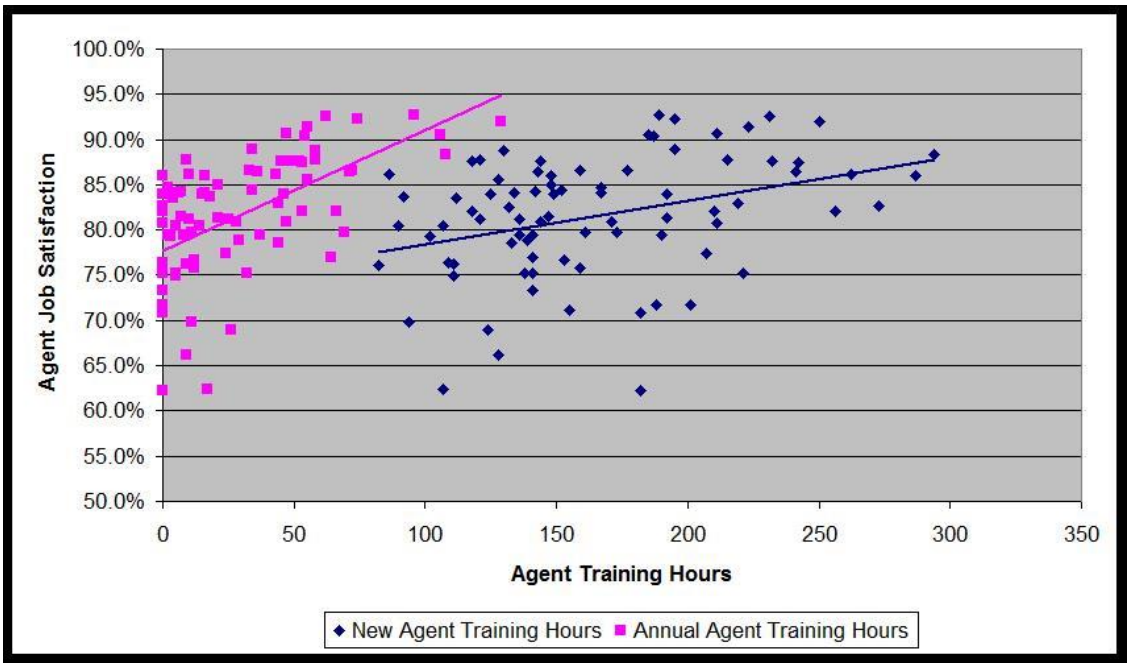

Industry Benchmarks for Agent Training

Agent training hours can vary dramatically from one organization to another, and from one industry to another. MetricNet's global benchmarking database shows that the average new agent training is 165 hours (roughly four full weeks), and the average annual agent training is 29 hours. These metrics range widely, however, from a low of 82 hours to a high of 294 hours for new agent training, and from a low of 0 hours to a high of 129 hours for annual agent training (Figure 3). These large variations are a reflection of each company's individual training philosophy, as well as budget and time constraints as mentioned above.

Figure 3: Industry Benchmarks for Agent Training Hours

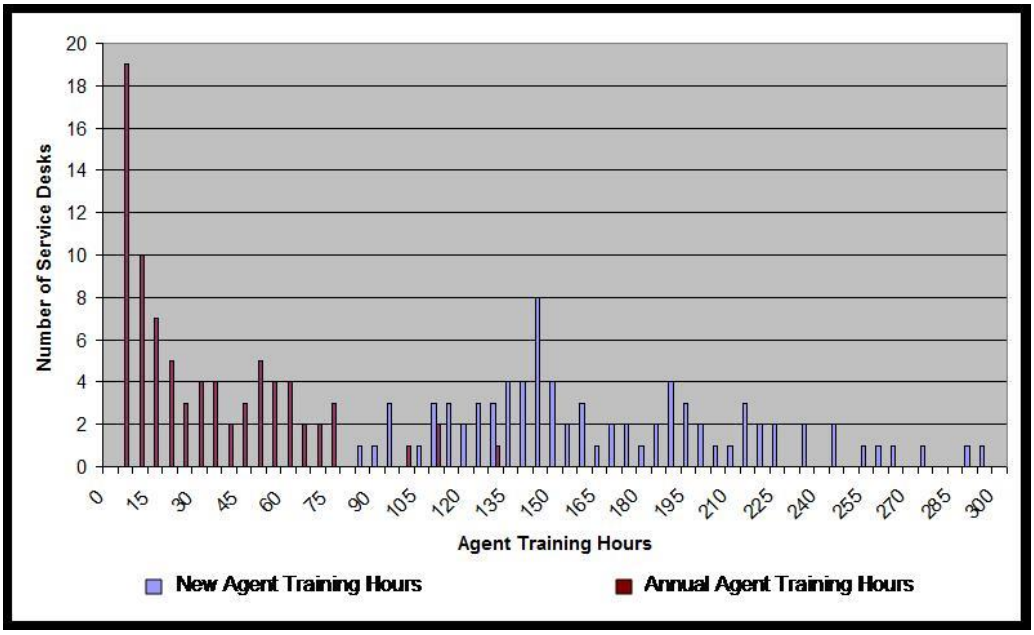

Cause-and-Effect for Service Desk KPIs

This month we depart from our usual format, and instead of discussing a single metric I will explore the cause-and-effect relationships between service desk KPIs.

Many of us have heard the sage advice "*You can't manage what you don't measure*." This is particularly true for the service desk, where effective performance measurement is not just a necessity, but a prerequisite for effective decision-making. Despite the widespread belief in this statement, few service desks use KPIs to their full potential. In fact the vast majority of service desks use metrics to track and trend their performance – but nothing more! Unfortunately, in this mode a service desk misses the real value of performance measurement by failing to exploit the *diagnostic* capabilities of KPIs. But the true potential of KPIs can only be unlocked when they are used holistically, not just to *measure* performance, but also to diagnose and understand the underlying drivers of performance.

The key to using KPIs diagnostically is to understand their cause-and-effect relationships. You can think of these relationships as a linkage where all of the KPIs are interconnected. When one KPI moves up or down, other KPIs invariably move with it. Understanding this linkage is enormously powerful because it provides insight into the levers you can pull to achieve desired outcomes.

The diagram in Figure 1 represents the service desk KPI linkage, and will be central to our discussion today. The metrics shown in red have been the subject of past Metric of the Month articles, which can be found in the archives of the Industry Insider.

Service Desk KPIs Figure 1: Service Desk Cause-and-Effect Diagram

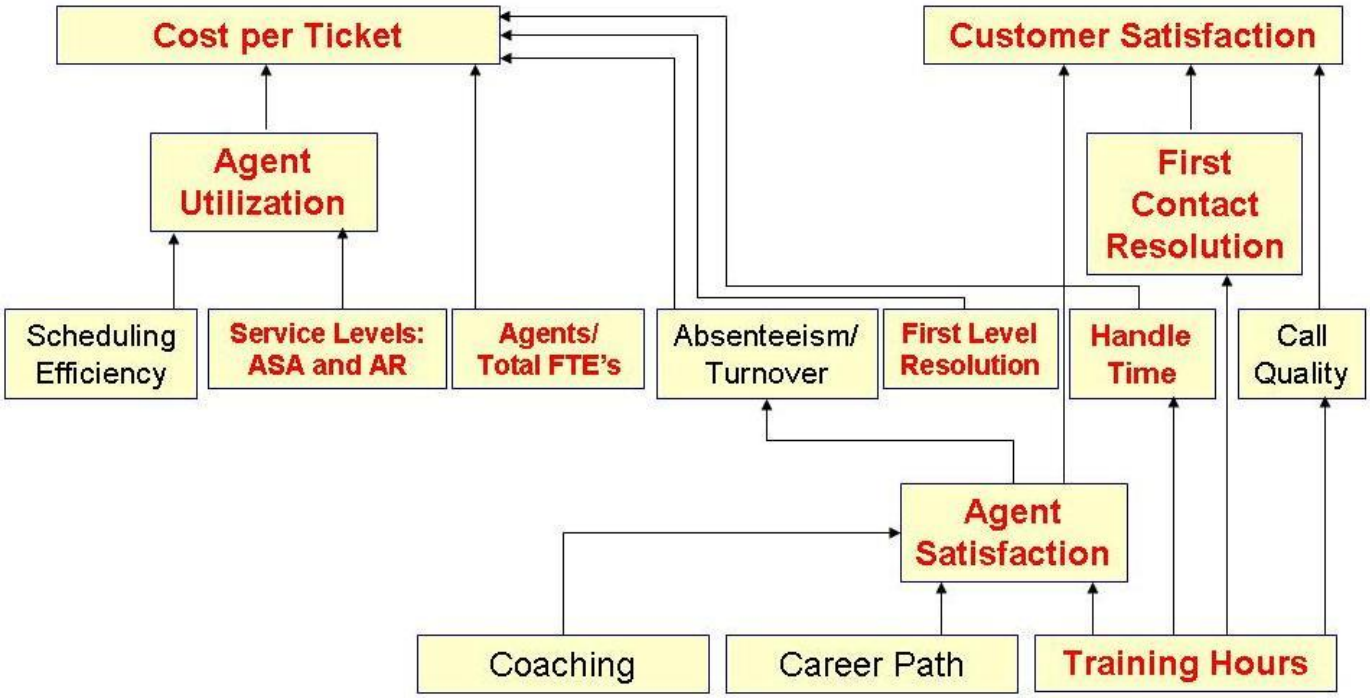

The Foundation Metrics

Virtually everything a service desk does can be viewed through the lens of cost and quality. Will this new technology reduce my costs? Will this new process improve customer satisfaction? This insight is crucial because it greatly simplifies decision-making on the service desk. Any undertaking that does not have the long-term effect of improving customer satisfaction, reducing costs, or both, is simply not worth doing. This is why cost per ticket and customer satisfaction are known as the foundation metrics.

These metrics are also very useful for telling the story of service desk performance. Most people instinctively understand cost and customer satisfaction, so it is easy to have a discussion about service desk performance in the context of these metrics. It is important to note, however, that the foundation metrics cannot be directly controlled. Instead, they are controlled through their underlying drivers.

The Underlying Drivers

Every KPI in the service desk is either directly or indirectly connected to cost per ticket and customer satisfaction. Those that directly impact the foundation metrics are called the underlying drivers, and include agent utilization, first contact resolution rate, and agent job satisfaction. Improvements in any of these metrics result in corresponding improvements in the foundation metrics. But unlike the foundation metrics, which cannot be directly controlled, the underlying drivers *can* be directly controlled. In fact this is where you have the greatest leverage to impact the cost and quality of your service desk.

If a service desk is struggling with high costs, for example, reducing the cost per ticket can often be achieved by increasing agent utilization or by reducing agent absenteeism and turnover. Likewise, if the goal of the service desk is to improve customer satisfaction, that can often be achieved by improving first contact resolution rate or call quality. The cause-and-effect relationship between first contact resolution rate and customer satisfaction has been discussed in prior Metric of the Month articles, and is shown in figure 2.

Service Desk KPIs Figure 2: First Contact Resolution Rate vs. Customer Satisfaction

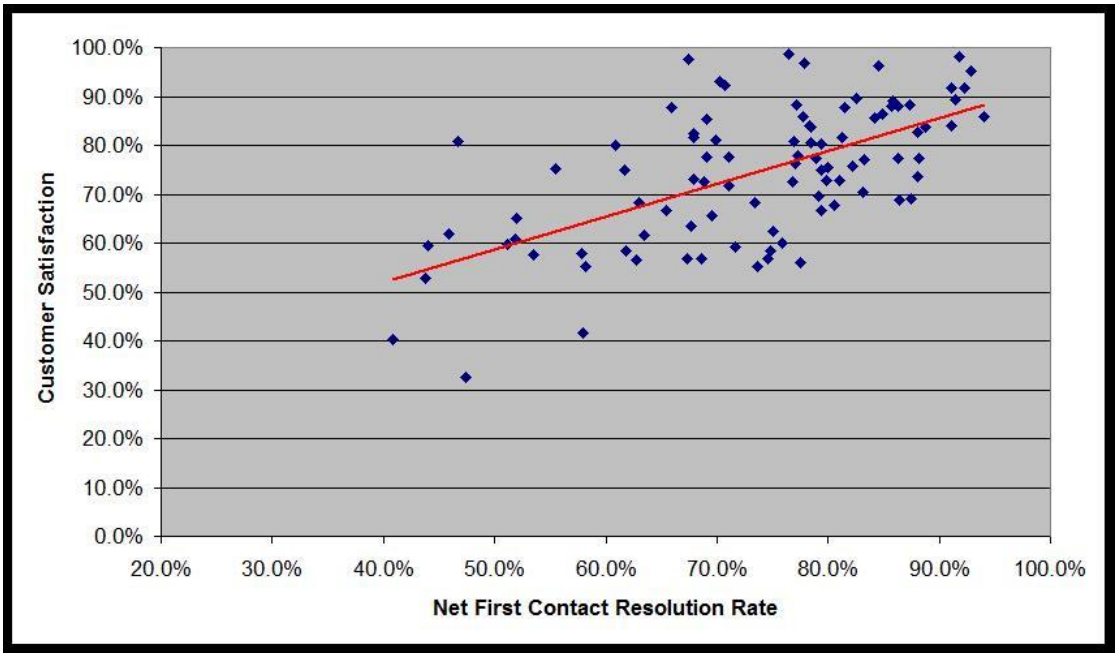

The Bellwether Metrics

Agent satisfaction and agent training hours are considered bellwether metrics because they are at the base of the KPI cause-and-effect diagram, and impact virtually every other metric in the service desk. Any movement in the bellwether metrics will be felt throughout the KPI linkage, and will eventually have an impact on the foundation metrics. If I know the agent satisfaction and training hours for a service desk, I can almost always predict what the cost and customer satisfaction will be.

High levels of agent job satisfaction translate into lower absenteeism and turnover, which then translates into lower cost. Likewise training hours that are above average almost always have the effect of producing higher first contact resolution rates and call quality, which then drives higher customer satisfaction levels. Moreover, training hours are also one of the key drivers of agent job satisfaction, and therefore represent a high leverage opportunity for a service desk to improve both its cost and quality performance. Figure 3 below shows the impact of training hours on agent job satisfaction.

Service Desk KPIs Figure 3: Agent Training Hours vs. Agent Job Satisfaction

Once you become familiar with the cause-and-effect relationships of service desk KPIs you will be in a much better position to identify, diagnose, and act upon any performance gaps in your service desk. This includes positive performance gaps, which you want to perpetuate, as well as negative performance gaps, which you can eliminate by modifying the underling drivers.

Tickets per Technician per Month

Tickets per Technician per Month is a Desktop Support metric. As the name suggests, Tickets per Technician per Month is simply the total number of monthly tickets logged by desktop support divided by the number of technicians working in desktop support.

For purposes of this discussion it is important to remember that desktop support tickets are comprised of both incidents and service requests. An Incident is typically unplanned work that requires the assistance of an on-site technician to resolve. Common examples include a desktop or laptop computer break/fix, a printer or server failure, connectivity problems, or any other issue that cannot be resolved remotely by the Level 1 Service Desk. By contrast, most Service Requests represent planned work. Among the most common Service Requests are Move's/Add's/Change's, hardware refresh/replacement, and device upgrades. Tickets represent the sum of all Incidents and Service Requests, as illustrated in Figure 1 below.

Figure 1: Tickets, Incidents and Service Requests

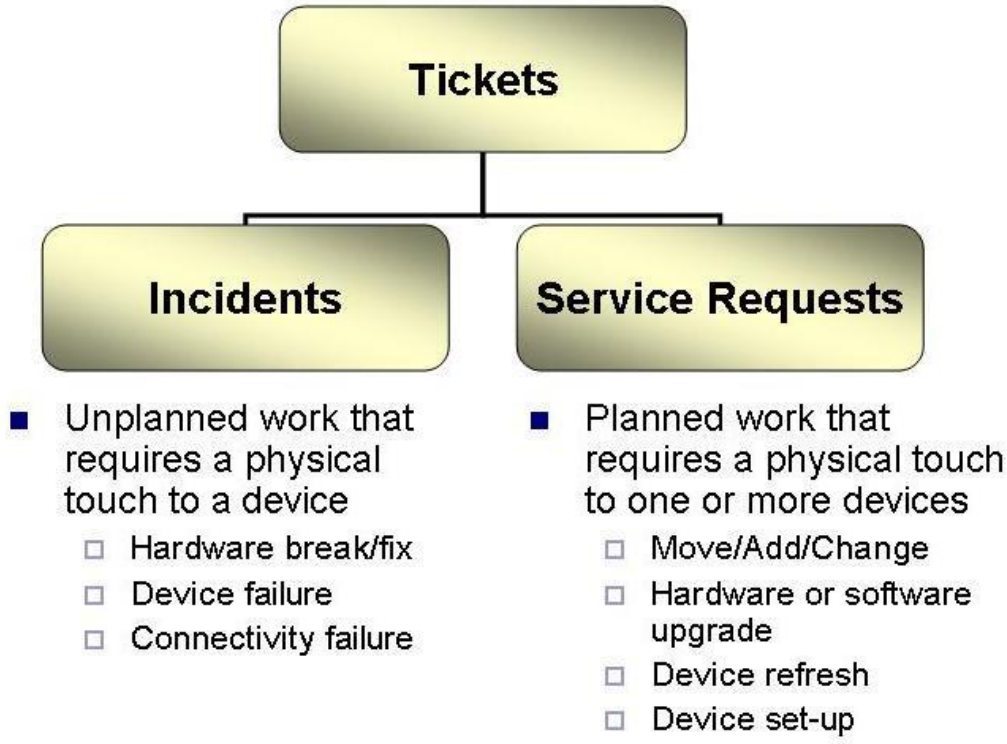

Incident Volume + Service Request Volume = Ticket Volume

Why it's Important

Tickets are the primary unit of work in desktop support. As such, ticket volume will drive the headcount of technicians needed by an organization. A common misperception in desktop support is that the user population alone will define the number of technicians needed. This approach wrongly assumes that the ratio of desktop support technicians to the number of users is fixed. For example, 12.5 desktop support technicians are needed for every 1,000 users. The error in this approach is that no two user populations have the same needs, and therefore no two user populations generate the same workload. As such, staffing decisions in desktop support should be based upon *workload*, not user population. With this in mind, it is easy to see why two organizations with exactly the same headcount may require very different staffing levels for desktop support.

Benchmark Data for Tickets per Technician per Month

The number of tickets that can be handled by a technician is driven by numerous factors including the mix of incidents and service requests, the average work time per ticket, and the average travel time per ticket. Service requests generally take more time to complete than incidents, so for organizations with a higher percentage of service requests vs. incidents the number of tickets that a technician can handle in a month will be lower. Likewise, as the work time per ticket increases, the number of tickets that a technician can handle in a month will decrease. And finally, in a high density user environment (think high rise office building) the travel time per ticket will be low, and this, in turn, increases the number of tickets a technician can handle in a month.

Figure 1 below shows just how dramatically the travel time, and work times for incidents and service requests can vary from one environment to another. These wide variations produce correspondingly wide ranges in the volume of tickets that a technician can handle in a month.

Figure 1: Desktop Support Work Time and Travel Time

Key Performance Indicator	Statistic	Financial Services	High Tech	Equipment Manufacturing	Telecom	Business Services	Health Care	Energy Utilities
Incident Work Time (minutes)	Average	18.3	19.8	14.7	16.1	21.5	12.3	14.2
	Min	6.3	12.8	5.1	7.4	12.1	6.6	6.1
	Max	71.7	65.9	44.0	73.5	58.0	59.0	30.3
Service Request Work Time (minutes)	Average	83.4	95.9	50.5	76.2	72.8	35.4	53.4
	Min	52.3	43.4	20.6	31.7	27.7	22.4	19.1
	Max	243	302	173	205	168	140	124
Travel Time per Ticket (minutes)	Average	25	32	33	19	27	20	12
	Min	11	17	14	9	11	8	7
	Max	110	153	164	64	79	52	53

Since ticket work and travel time drives technician headcount, one would also expect to see a wide variation in headcount requirements from company to company, and from industry to industry. Figures 2 below shows that the average desktop technician headcount can range from a low of just 5.4 technicians per 1,000 seats in health care, to a high of 28.4 technicians per 1,000 seats in the high tech industry.

Figure 2: Desktop Support Technician Headcount Ranges by Industry

Key Performance Indicator	Statistic	Financial Services	High Tech	Equipment Manufacturing	Telecom	Business Services	Health Care	Energy Utilities
Desktop Technicians per 1,000 Seats Supported	Average	21.9	28.4	12.7	15.5	27.0	5.4	7.7
	Min	3.4	7.5	2.4	2.6	3.6	1.1	1.2
	Max	67.3	90.1	40.5	40.4	55.2	12.3	16.3

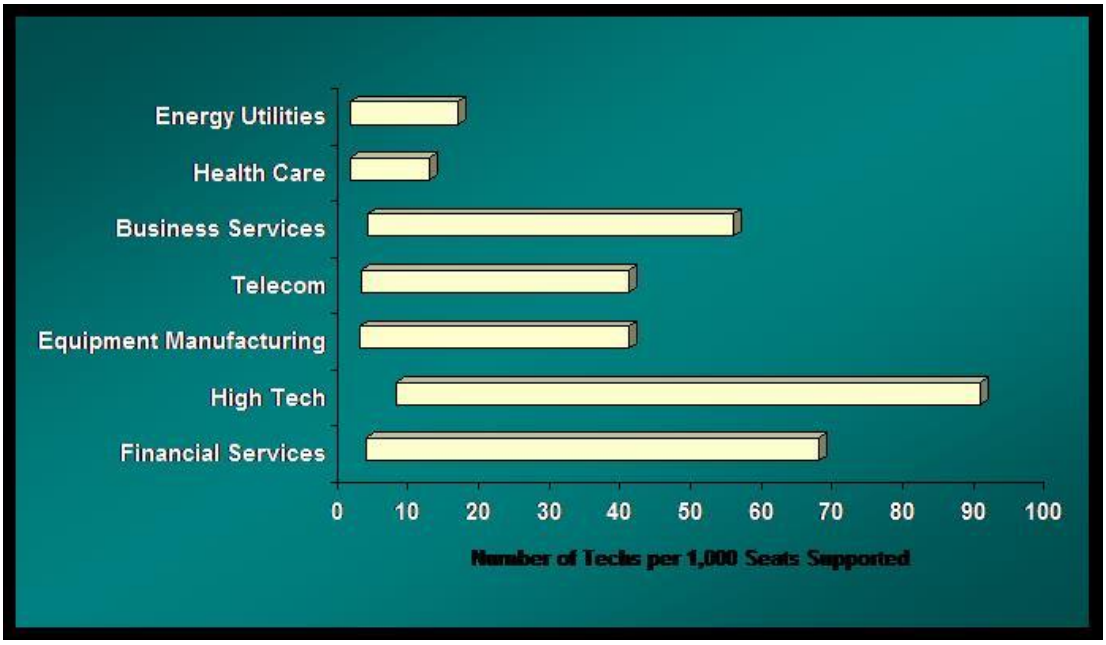

Finally, when we look at the benchmarking data for the number of Tickets per Technician per month in Figure 3, we see a wide range of values, from a low of 30 Tickets per Technician per Month, to a high of 198 Tickets per Technician per Month. These wide variations are once again driven by differences in travel time and work time per ticket, which are in turn a function of the unique environment in which each technician operates.

Figure 3: Tickets per Technician per Month by Industry

Key Performance Indicator	Statistic	Financial Services	High Tech	Equipment Manufacturing	Telecom	Business Services	Health Care	Energy Utilities
Tickets per Technician per Month	Average	99	100	98	124	87	108	133
	Min	43	58	30	58	48	78	67
	Max	191	136	147	168	123	198	167

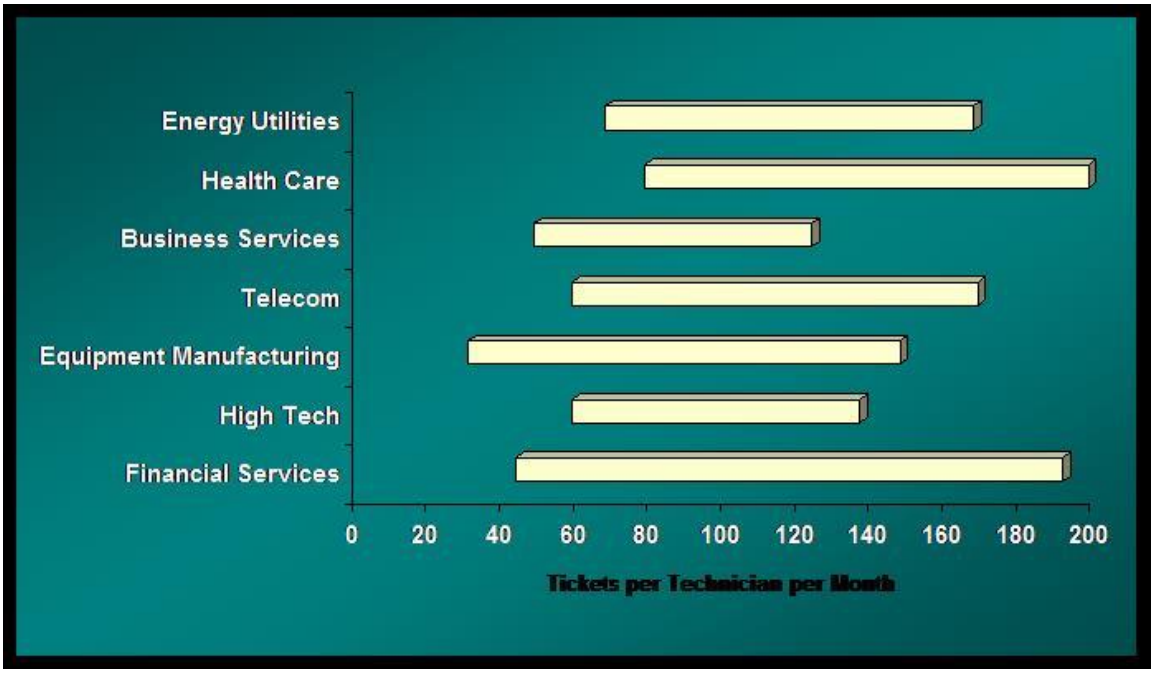

Annual Agent Turnover

Annual Agent Turnover is the percentage of all agents that leave a support organization over the course of a year. Let's say, for example, that the average agent headcount for your service desk is 15, and that 5 agents leave and must to be replaced during the year. The annual agent turnover would therefore be 5 ÷ 15 = 33.3%. The metric is equally applicable to the service desk and desktop support, but the examples and illustrations we use in this month's article will be specific to the service desk.

Some organizations make a distinction between good turnover and bad turnover. Bad turnover is when an agent leaves the company altogether because of performance issues or to pursue other job opportunities. So-called good turnover, by contrast, is when an agent who is otherwise performing well is moved or promoted to a non-customer facing position in the service desk, or accepts another position in the company that is outside of the service desk. Both types of turnover are included in the calculation of annual agent turnover because both types of turnover create a vacancy that must to be filled.

Why it's Important

Agent turnover can be detrimental to a service desk because it typically results in a seasoned agent being replaced by a less experienced agent. When there is turnover, the knowledge and experience of the agent leaving the service desk walks out the door with them. For those who have worked in a service desk, you know how painful this can be! Industry estimates place the cost of replacing an agent at more than $10,000 in North America. This includes the cost of identifying, screening, recruiting, and training a new agent, as well as the indirect cost of lower productivity that results when a new agent encounters the learning curve of a new job.

It has been shown in a past Metric of the Month article that one of the primary cause-and-effect relationships in the service desk is between agent job satisfaction and agent turnover: high job satisfaction is strongly correlated with low agent turnover rates and vice versa. The reasons for this are fairly obvious. When agents are satisfied with their work life, they tend to say put. When they are unhappy at work, they are more likely to leave. It is important to note, however, that turnover *can* be controlled by proactively managing agent job satisfaction. Once again, in a past Metric of the Month article, we have shown that effective career planning, training, and coaching can drive high levels of agent job satisfaction, and reduce turnover accordingly.

Benchmark Data for Annual Agent Turnover

It is common knowledge that technical service and support is a high turnover industry. Some of this turnover is due to the high stress nature of the job. While job stress may be a non-controllable factor, certain other factors, as mentioned above, can be managed and controlled to reduce and minimize turnover.

As shown in Figure 1, the average turnover in the industry is nearly 40% per year! This means that the average service desk agent stays for just 2 ½ years before moving on.

Figure 1: Annual Agent Turnover

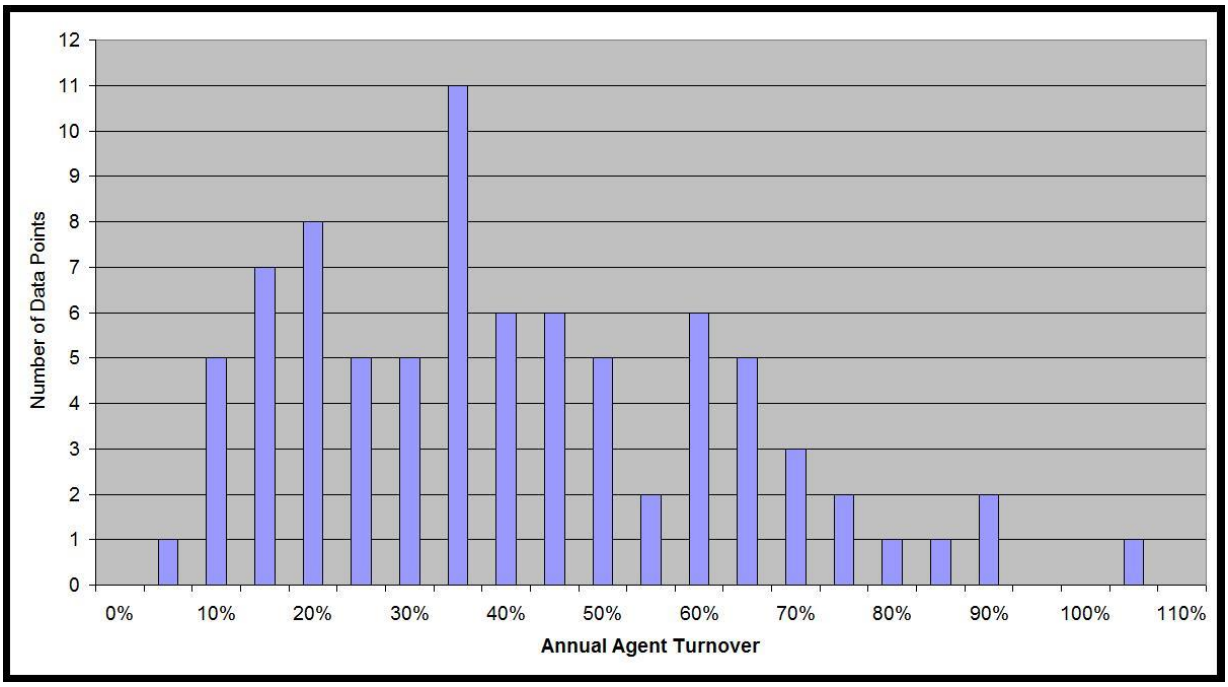

Figures 2 and 3 show how agent turnover affects some of the most important KPI's in the service desk. Since agent turnover is inversely related to agent experience levels, one can deduce from these benchmarks that less experienced agents have a negative impact on both first contact resolution rate and customer satisfaction. In general, the more experienced the agent pool, the higher the first contact resolution rate and customer satisfaction will be. So it turns out that agent turnover is not only costly, it also has an adverse impact on service desk performance!

Figure 2: Agent Turnover vs. First Contact Resolution Rate

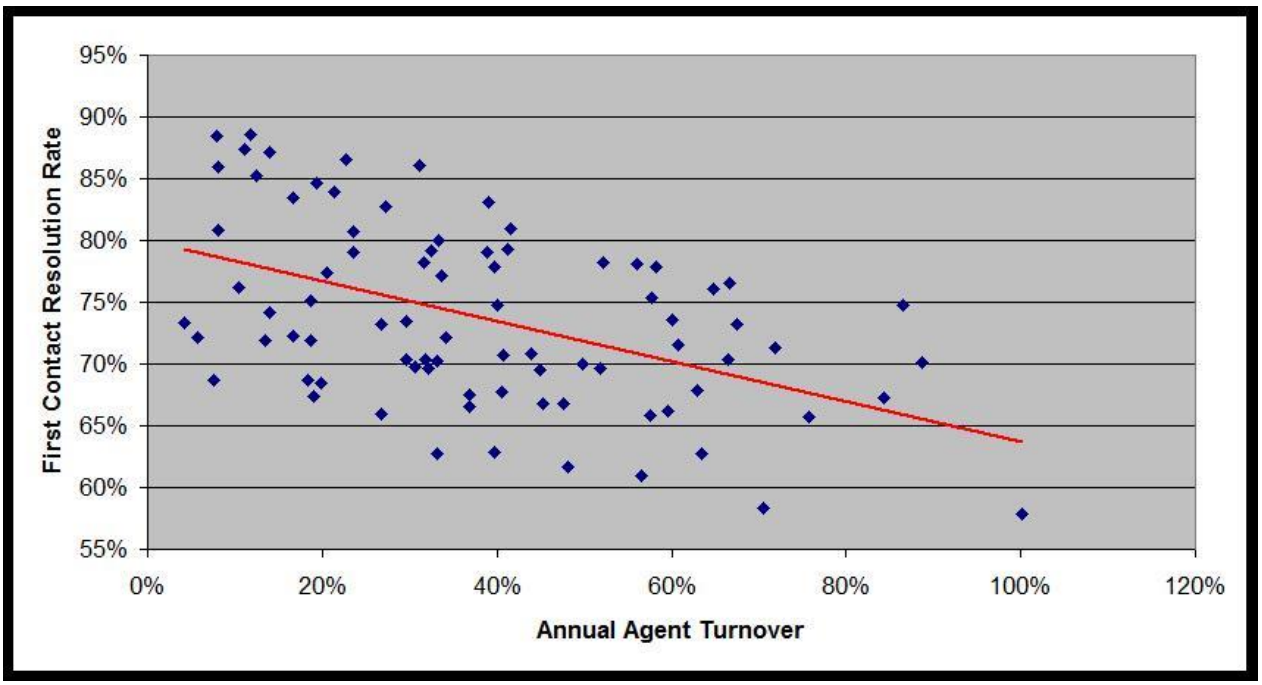

Figure 3: Agent Turnover vs. Customer Satisfaction

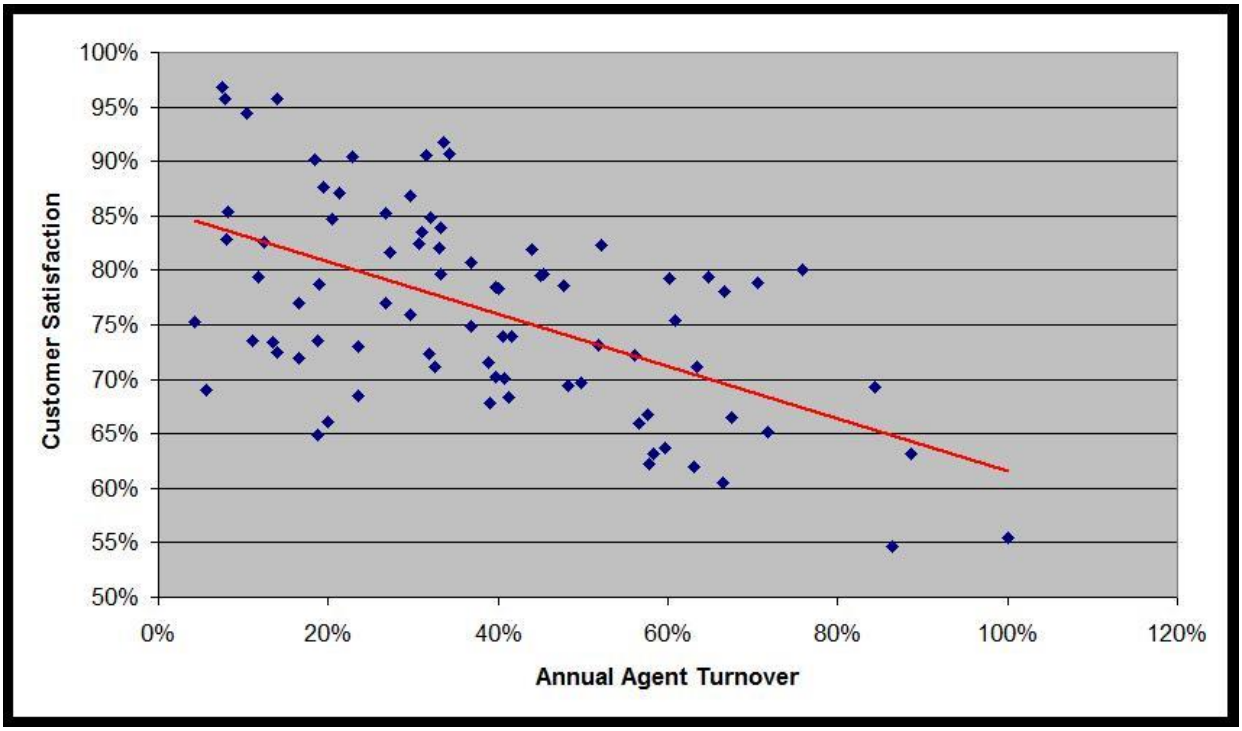

Cause-and-Effect for Desktop Support KPI

This month we depart from our usual format, and instead of discussing a single metric I will explore the cause-and-effect relationships between desktop support KPIs. A companion article to this one, which explored the cause-and-effect relationship of service desk KPIs, was published in the November 2012 issue of Industry Insider.

Cause-and-Effect for Desktop Support KPIs

Many of us have heard the sage advice "*You can't manage what you don't measure*."

This is particularly true for desktop support, where effective performance measurement is not just a necessity, but a prerequisite for effective decision-making. Despite the widespread belief in this statement, few desktop support groups use KPIs to their full potential. In fact the vast majority of desktop support groups use metrics to track and trend their performance – but nothing more! Unfortunately, in this mode desktop support misses the real value of performance measurement by failing to exploit the *diagnostic* capabilities of KPIs. But the true potential of KPIs can only be unlocked when they are used holistically, not just to *measure* performance, but also to diagnose and understand the underlying drivers of performance.

The key to using KPIs diagnostically is to understand their cause-and-effect relationships. You can think of these relationships as a linkage where all of the KPIs are interconnected. When one KPI moves up or down, other KPIs invariably move with it. Understanding this linkage is enormously powerful because it provides insight into the levers you can pull to achieve desired outcomes.

The diagram in Figure 1 represents the desktop support KPI linkage, and will be central to our discussion today. The metrics shown in red have been the subject of past Metric of the Month articles, which can be found in the archives of the Industry Insider.

Figure 1: Desktop Support Cause-and-Effect Diagram

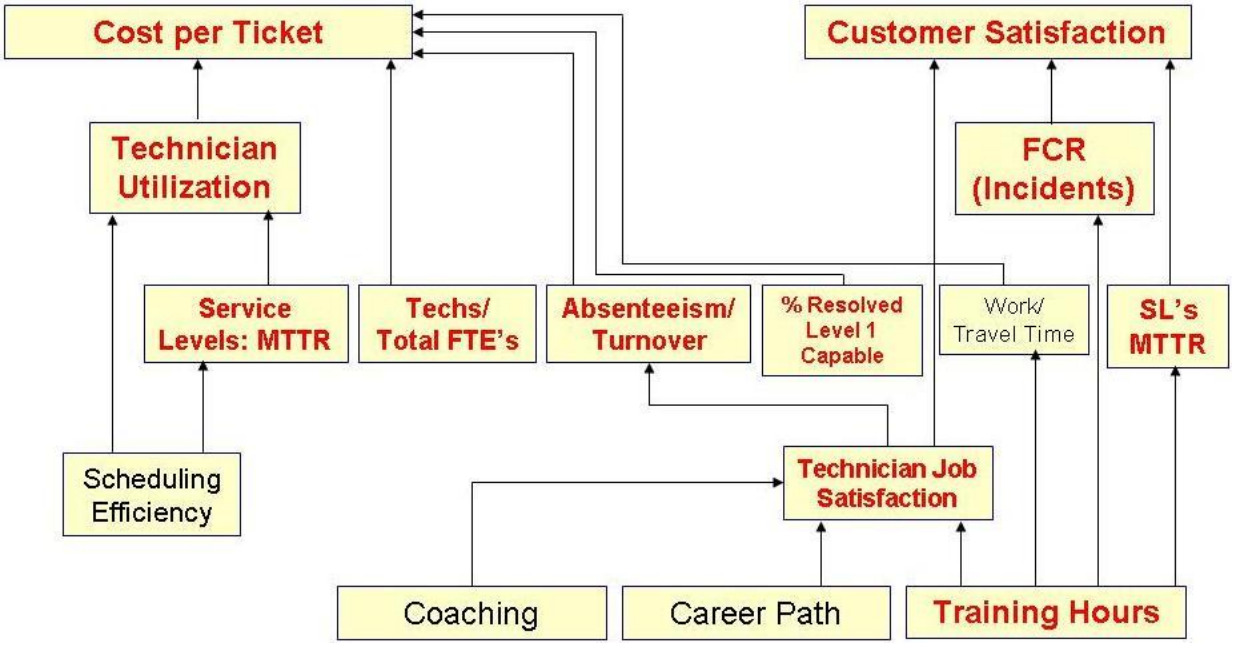

The Foundation Metrics

Virtually everything undertaken by desktop support can be viewed through the lens of cost and quality. Will this new technology reduce my costs? Will this new process improve customer satisfaction? This insight is crucial because it greatly simplifies decision-making for desktop support. Any undertaking that does not have the long-term effect of improving customer satisfaction, reducing costs, or both, is simply not worth doing. This is why cost per ticket and customer satisfaction are known as the foundation metrics.

These metrics are also very helpful for telling the story of desktop support performance. Most people instinctively understand cost and customer satisfaction, so it is easy to have a discussion about the performance of desktop support in the context of these two metrics. It is important to note, however, that the foundation metrics cannot be directly controlled. Instead, they are controlled through their underlying drivers.

The Underlying Drivers

Every KPI in desktop support is either directly or indirectly connected to cost per ticket and customer satisfaction. Those that directly impact the foundation metrics are called the underlying drivers, and include technician utilization, the ratio of technicians to total headcount, first contact resolution rate for incidents, technician job satisfaction, and mean time to resolve (MTTR). Improvements in any of these metrics result in corresponding improvements to the foundation metrics. But unlike the foundation metrics, which cannot be directly controlled, the underlying drivers *can* be directly controlled. In fact this is where you have the greatest leverage to impact the cost and quality of desktop support.

If a desktop support group is struggling with high costs, for example, reducing the cost per ticket can often be achieved by increasing technician utilization or by reducing technician absenteeism

and turnover. Likewise, if the goal of desktop support is to improve customer satisfaction, this can often be achieved by improving the main service level metric, mean time to resolve. The cause-and-effect relationship between incident MTTR and customer satisfaction has been discussed in a prior Metric of the Month article, and is shown in figure 2.

Figure 2: Incident Mean Time to Resolve vs. Customer Satisfaction

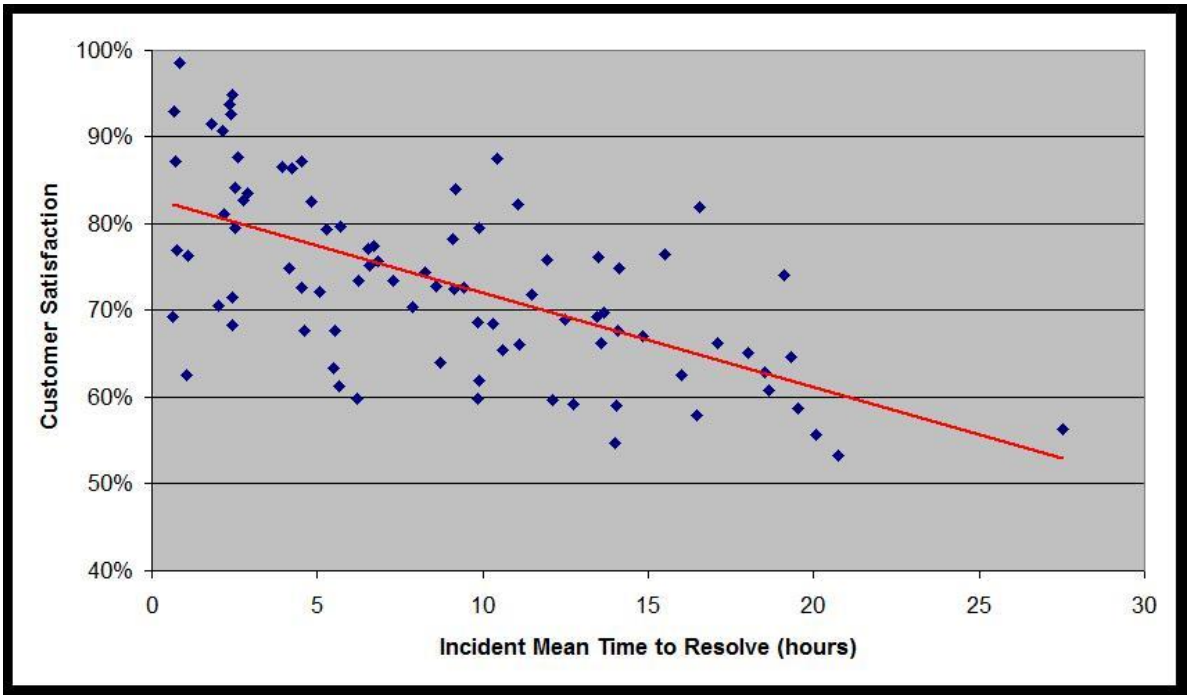

The Bellwether Metrics

Technician job satisfaction and technician training hours are considered bellwether metrics because they are at the base of the KPI cause-and-effect diagram, and impact virtually every other metric in desktop support. Any movement in the bellwether metrics will be felt throughout the KPI linkage, and will eventually have an impact on the foundation metrics. If I know the technician satisfaction and training hours for a desktop support group, I can almost always predict what the cost and customer satisfaction will be.

High levels of technician job satisfaction translate into lower absenteeism and turnover, which then translates into lower cost. Likewise training hours that are above average almost always have the effect of producing a higher first contact resolution rate for incidents, which then drives higher customer satisfaction levels. Moreover, training hours are also one of the key drivers of technician job satisfaction, and therefore represent a high leverage opportunity for desktop support to improve both its cost and quality performance. Figure 3 below shows the impact of training hours on technician job satisfaction.

Figure 3: Technician Training Hours vs. Technician Job Satisfaction

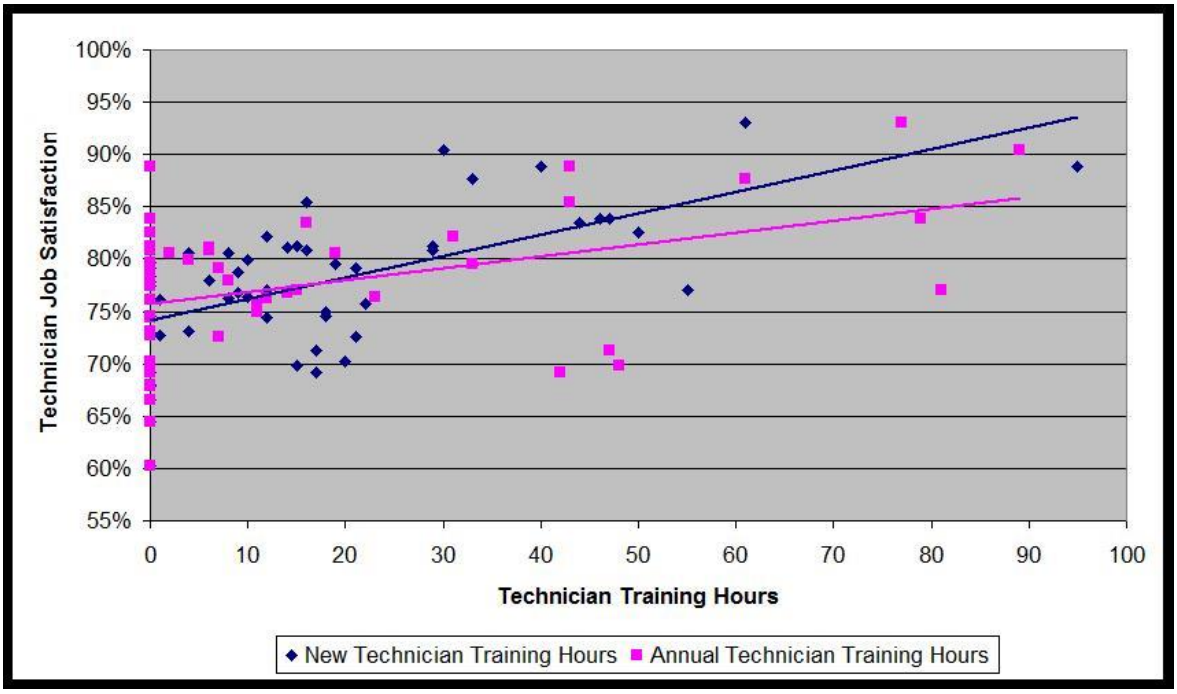

Once you become familiar with the cause-and-effect relationships of desktop support KPIs you will be in a much better position to identify, diagnose, and act upon any performance gaps in desktop support. This includes positive performance gaps, which you want to perpetuate, as well as negative performance gaps, which you can eliminate by modifying the underling drivers.

Net Promoter Score

Net Promoter Score, or NPS, is based on the idea that every organization's customers can be divided into three categories: Promoters, Passives, and Detractors. By asking one question — How likely is it that you would recommend our service to a friend or colleague? — you can track these groups and get a clear measure of your support organization's performance from the customer's perspective. Customers respond on a 0-to-10 point rating scale and are categorized as follows:

- **Promoters** (score 9-10) are loyal enthusiasts who will refer others to your support organization.
- **Passives** (score 7-8) are satisfied but unenthusiastic customers who may choose another source of support if given the chance.
- **Detractors** (score 0-6) are unhappy customers who can damage your reputation through negative word-of-mouth.

To calculate your support group's NPS, simply take the percentage of customers who are Promoters and subtract the percentage who are Detractors. Your Net Promoter score can be as low as –100% (everybody is a detractor) or as high as +100% (everybody is a promoter). An NPS that is positive (i.e., higher than zero) is thought to be good, and an NPS of +50% or greater is excellent.

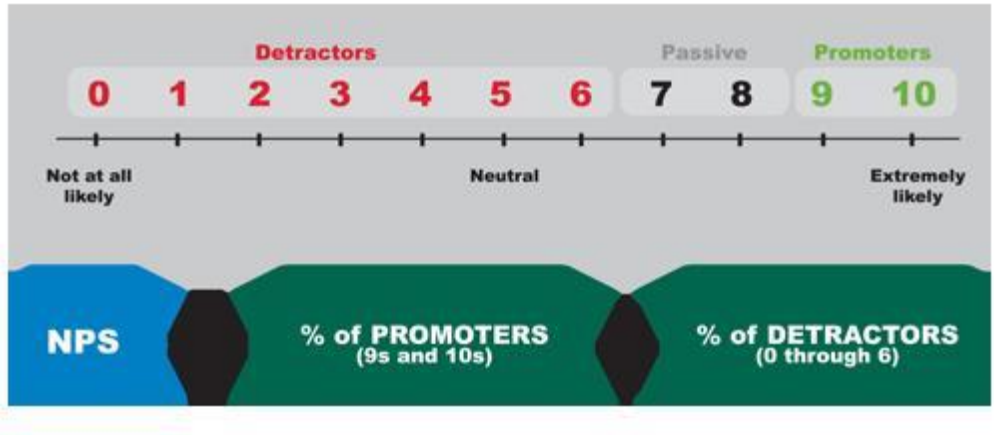

IT support groups that track NPS will typically follow the initial question with an open-ended request for elaboration, soliciting the reasons for a customer's rating of their IT support. These reasons can then be provided to front-line employees and management teams for follow-up action and improvement initiatives.

Benchmark Data for NPS

NPS is a relatively new metric, and is currently tracked by approximately 20% of all service desks and desktop support groups. However, the vast majority of organizations that track NPS are providing support to external businesses or consumers, where NPS is considered to be a leading indicator of customer loyalty and follow-on business.

MetricNet's benchmarking data for Customer Satisfaction and Net Promoter Scores is shown in Figure 1 for 68 IT support organizations that track both metrics. As you can see, the Customer Satisfaction score for this particular data set is always higher than the Net Promoter score. This is generally the case, as it is extremely rare for an organization to have a higher NPS than Customer Satisfaction score. Put another way, it is much easier to create a satisfied customer than it is to create a customer who is willing to put their credibility on the line by recommending your service.

Figure 1: Customer Satisfaction and Net Promoter Score

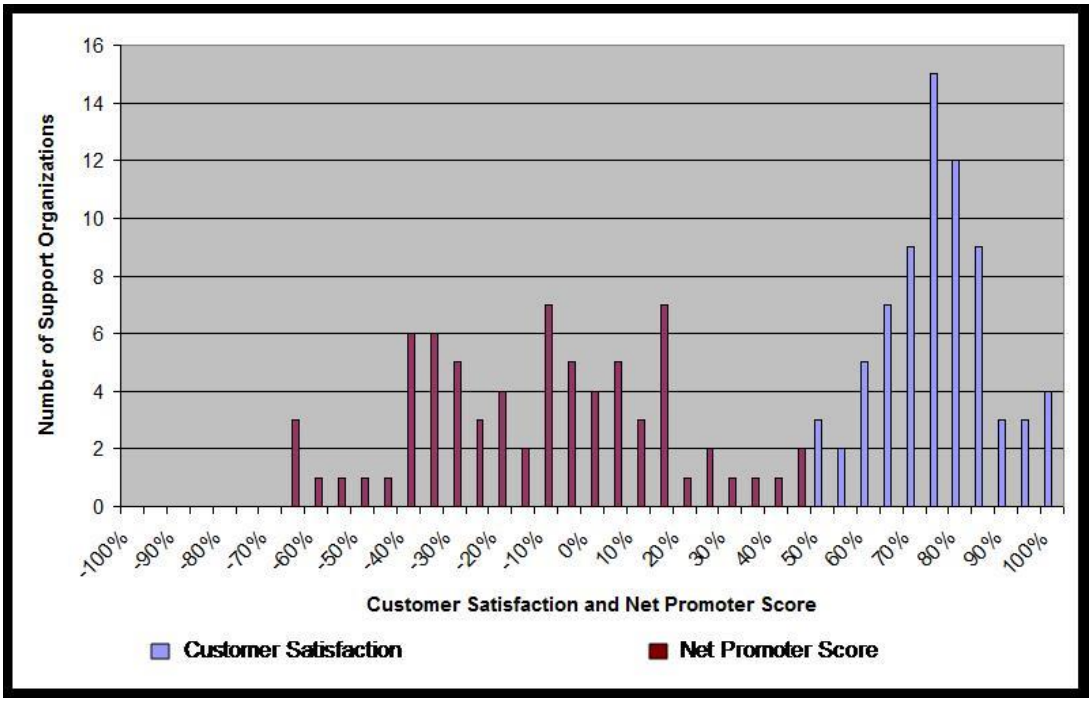

Since NPS is both a quality and a loyalty metric, one would expect there to be a correlation between NPS and Customer Satisfaction. There is, in fact, a correlation between these two metrics, but it is not as strong as one might expect. Figure 2 shows the strength of this correlation for the same 68 support organizations referenced in Figure 1. This reinforces the notion that NPS measures something different than customer satisfaction (i.e., customer loyalty), and that customer loyalty and satisfaction are two very distinct metrics.

Figure 2: Net Promoter vs. Customer Satisfaction

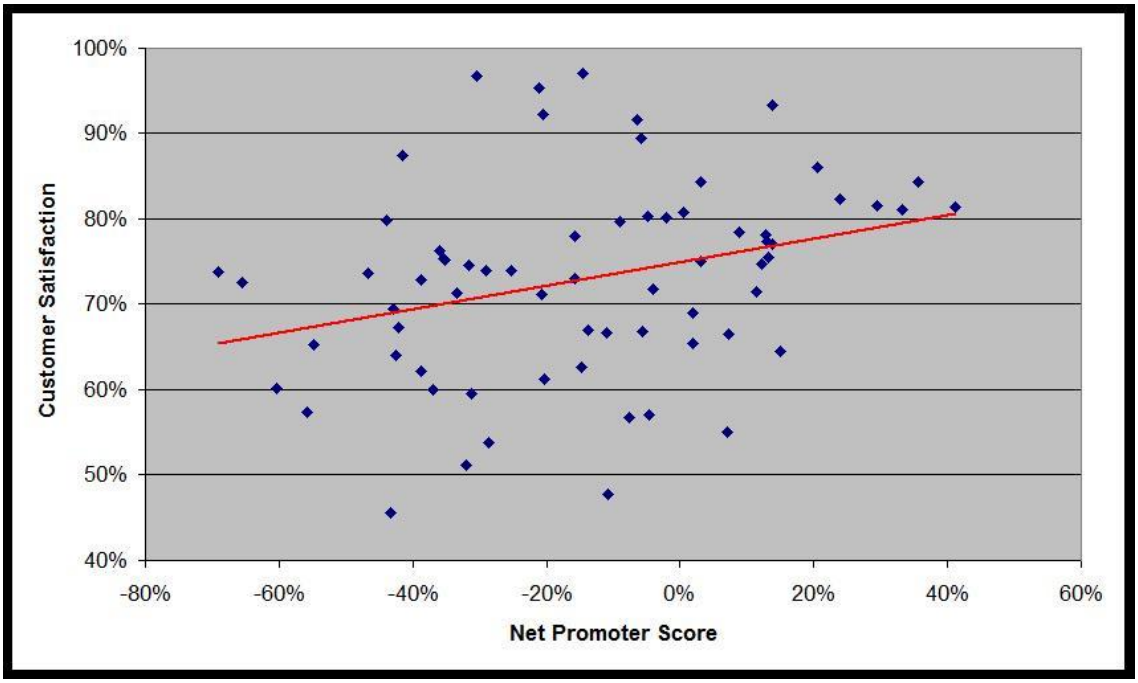

The key question with any IT service and support metric is whether it helps the organization operate more efficiently and effectively. Does it provide useful feedback that helps agents and technicians deliver exceptional customer service? Does it provide a tracking mechanism that enables week-over-week, and month-over-month comparisons? Those who track this metric believe that it does. But it is worth noting once again that NPS is used primarily by IT service and support organizations that are supporting an external customer base.

Service Desk Balanced Scorecard

This month we depart from our usual format, and instead of discussing a single metric I will explain how a handful of critical metrics can be combined to create a single, overall indicator of service desk performance. We call this the service desk balanced scorecard.

The Balanced Scorecard

Today's service desk technologies and reporting packages make it easy to capture copious amounts of performance data. Most service desk managers can tell you everything from last month's average speed of answer to yesterday's average handle time. But what does it all mean? If my abandonment rate goes up, but my cost per contact goes down, is that good or bad? Is my service desk performing better this month than it was last month?

Despite all the data that service desk managers have at their fingertips, most cannot answer a very basic question: How is my service desk performing? The balanced scorecard resolves this dilemma by combining the most important service desk KPI's into a single, overall measure of service desk performance.

MetricNet's research shows that establishing an overall metric for your service desk is critical. We call this metric the balanced score because it truly does communicate a balanced picture of service desk performance. The balanced scorecard is a mechanism that aggregates the most important service desk metrics – such as cost per contact and customer satisfaction – into a single, all-inclusive measure of service desk performance.

The value of this metric, when tracked over time, is that it enables a service desk to determine whether overall performance is improving or declining.

Oftentimes, when a service desk attempts to communicate its performance to other stakeholders in the business, particularly to lay people who do not understand service desk operations, those people quickly become overwhelmed by the minutia of such measures as first contact resolution rate and speed of answer, and they are confused about how to interpret the results. They are likely to focus in on one, easily-understood metric like speed of answer, and draw conclusions about overall service desk performance from this relatively unimportant metric. This is a classic case of missing the forest for the trees. It is therefore absolutely critical to communicate the overall performance of the service desk, and the balanced scorecard does that for you. Think of the balanced scorecard as your letter grade for the month! In this way, a service desk can track its overall performance, and, in any given month, may see costs go up or customer satisfaction go down or speed of answer increase, but these individual measures take on a secondary level of importance because the balanced score provides a more complete and accurate picture of overall service desk performance.

The Mechanics of Creating a Scorecard

Creating a scorecard is relatively straightforward. You can follow along in Figure 1 below as I explain the process.

Figure 1: Service Desk Balanced Scorecard

Performance Metric	Metric Weighting	Performance Range		Your Performance	Metric Score	Balanced Score
		Worst Case	Best Case			
Cost per Inbound Contact	25.0%	$20.50	$7.09	$8.42	90.1%	22.5%
Customer Satisfaction	25.0%	63.1%	96.7%	87.6%	72.9%	18.2%
Technician Utilization	15.0%	30.8%	62.7%	54.9%	75.5%	11.3%
First Contact Resolution Rate	15.0%	57.6%	91.4%	57.6%	0.0%	0.0%
Technician Job Satisfaction	10.0%	53.3%	99.0%	99.0%	100.0%	10.0%
% of Calls Answered in 30 Seconds	10.0%	2.6%	85.8%	2.6%	0.0%	0.0%
Total	100.0%	N/A	N/A	N/A	N/A	62.1%

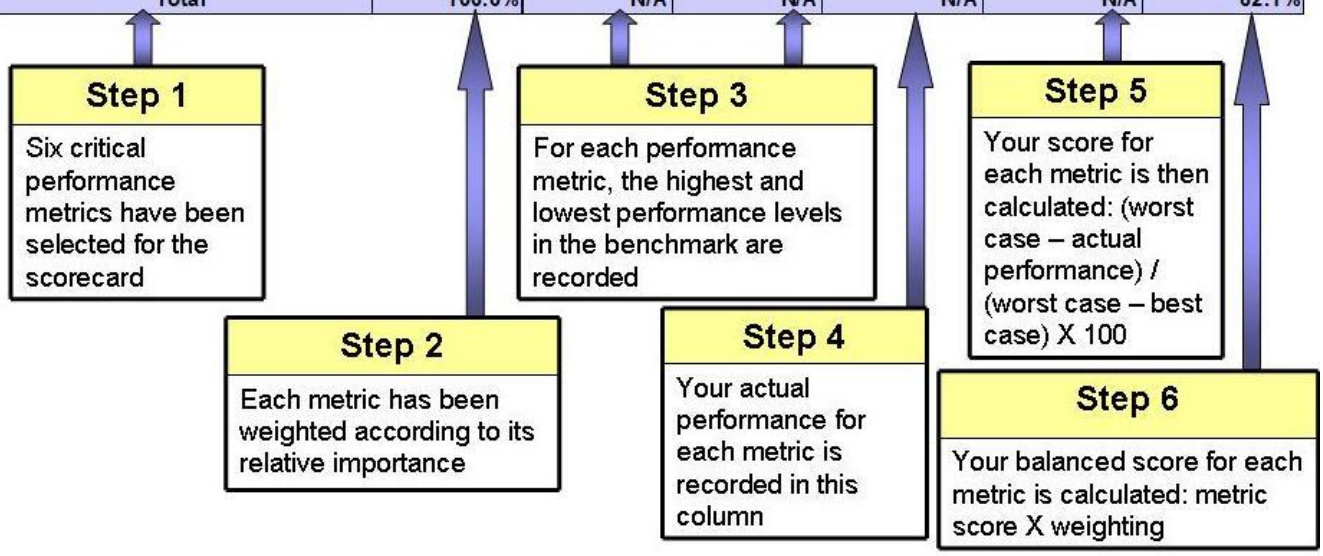

First you select the metrics to include in your scorecard. We suggest including the following six metrics: Cost per contact, customer satisfaction, agent utilization, first contact resolution rate, agent job satisfaction, and average speed of answer. Depending upon the metrics you track in your service desk, you may choose fewer metrics or a different mix of metrics for your scorecard. Secondly, you establish a weighting for each metric based upon its relative importance in the scorecard. This is a judgment call, but we suggest overweighting cost and customer satisfaction, since these are the foundation metrics for service and support. Step 3 is to show a reasonable range of performance – worst case to best case – for each metric. Normally these performance ranges are derived from a benchmark of your service desk. In step 4 your performance for each metric is inserted into the third column from the right. A score for each metric is then calculated based on the interpolation formula in step 5. And finally, a balanced score for each metric is determined by multiplying the metric weighting by the metric score. When the metric scores are summed up, you have the total balanced score for your service desk.

In this particular example, the service desk balanced score is 62.1%. Your balanced score will always range from 0% (if you have the worst possible performance for every metric in the

scorecard) to 100% (if you have the best possible performance for every metric in the scorecard). It turns out that the service desk in our example has scored quite well. When we run hundreds of service desks through this algorithm, we get a normal distribution centered right at 50%. Those who score above 61% are in the top quartile; those who score between 50% and 61 % are in the second quartile; those between 39% and 50% are in the third quartile; and those below 39% are in the bottom quartile for overall performance.

Benchmarking Your Performance

The balanced scorecard is an ideal way to track, trend, and benchmark your service desk performance. Figure 2 below shows the trend in one company's service desk performance over a 12 month period. The blue bars in the chart represent the monthly balanced scores, while the red background represents the 12 month trailing trend in scorecard performance. Clearly, the performance trend for this particular service desk is improving!

Figure 2: Balanced Scorecard Trend

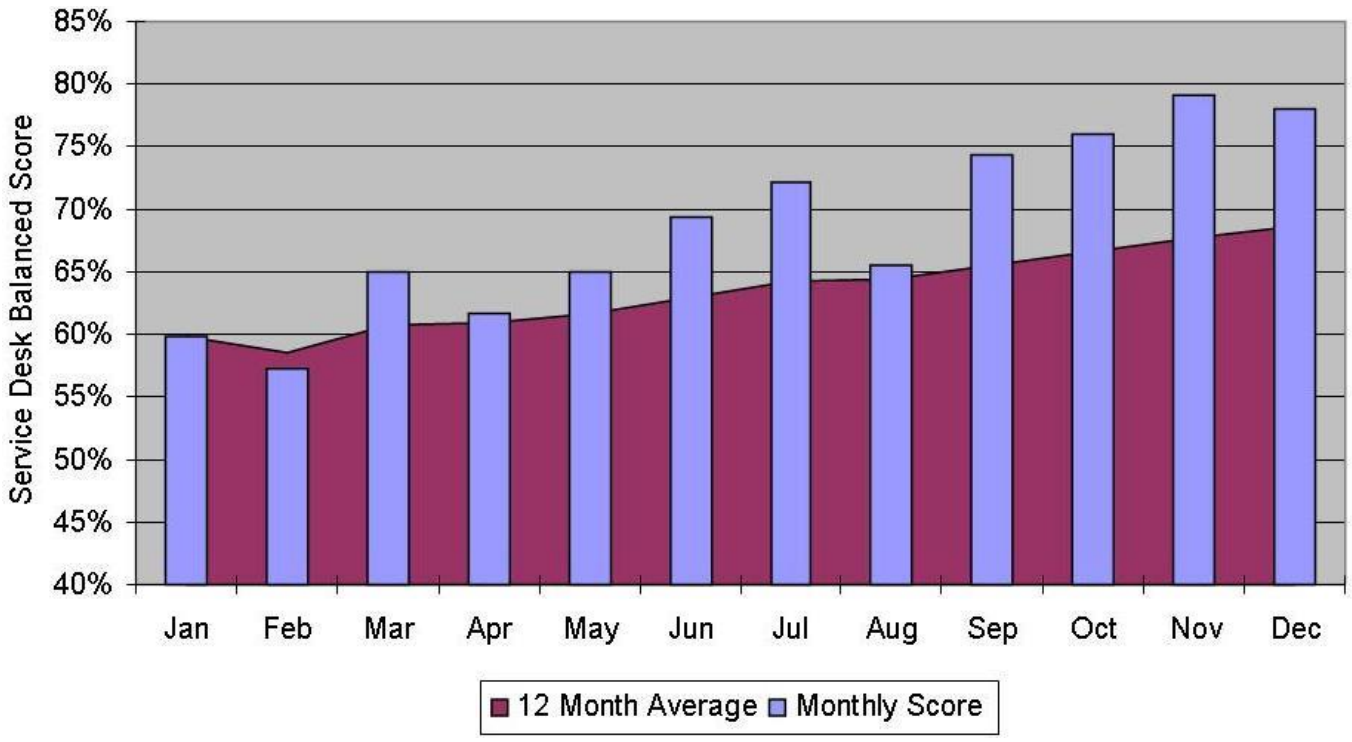

Finally, the service desk balanced score can be used to benchmark your service desk on a fair, apples-to-apples basis against other service desks. Figure 3 below shows how the service desk in our example compares to 23 other service desks in their benchmarking peer group.

Figure 3: Scorecard Benchmarking Comparison

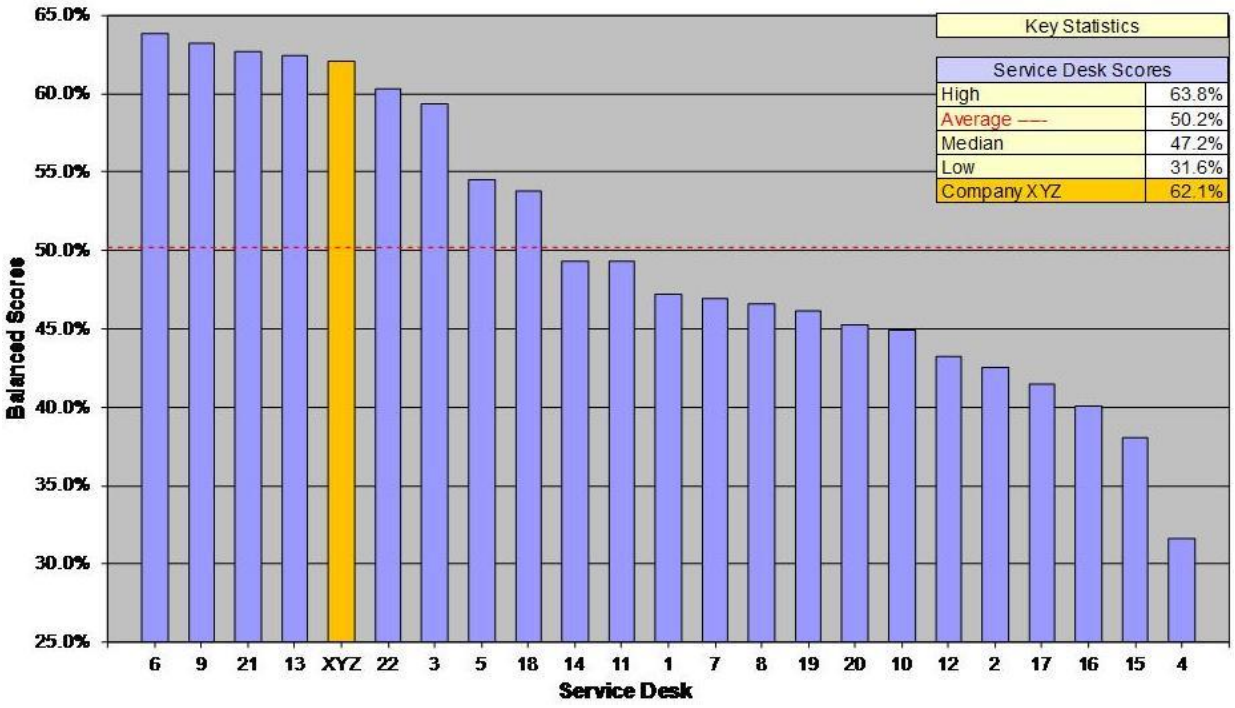
Key Statistics
Service Desk Scores
High 63.8%
Average ---- 50.2%
Median 47.2%
Low 31.6%
Company XYZ 62.1%
65.0%
60.0%
55.0%
50.0%
45.0%
40.0%
35.0%
30.0%
25.0%
Balanced Scores
6 9 21 13 XYZ 22 3 5 18 14 11 1 7 8 19 20 10 12 2 17 16 15 4
Service Desk

Service Desk Agent Occupancy

Agent occupancy is a service desk metric that measures the percentage of logged in time that an agent is actually delivering customer service. Let's say, for example, that an agent is logged into the service desk ACD for six hours a day, and is in talk, hold, or wrap mode for four of those six hours. In this case, agent occupancy would be calculated as 4 hours ÷ 6 hours = 66.7% agent occupancy. Likewise, let's say that an agent is logged into the ticketing system for seven hours a day, and is responding to customer tickets that are emailed to the service desk for five of those seven hours. The calculation for agent occupancy in this case would be 5 hours ÷ 7 hours = 71.4% agent occupancy.

Agent occupancy is oftentimes confused with agent utilization. Although the numerator is the same for both metrics, the denominator is different. The denominator for occupancy, as explained above, is the total time that a voice, email, or chat agent is logged into the system. By contrast, the denominator for agent utilization is the total time that a voice, email, or chat agent is at work, *including the time that the agent is logged into the system.*

To continue with our example from above, let's assume in both cases that the agent is at work for a total of eight hours. In the first example, the agent utilization would be 4 hours ÷ 8 hours = 50% agent utilization. Likewise, in our second example the agent utilization would be 5 hours ÷ 8 hours = 62.5% agent utilization. Agent utilization will always be less than or equal to agent occupancy.

While agent occupancy and utilization have a weak correlation (Figure 1 below), it is possible to have a high agent occupancy and a low agent utilization. This scenario occurs when agents are not logged into the system for very many hours each day, but are fairly busy when they *are* logged into the system. For this reason, agent utilization is recognized as a more accurate indicator than agent occupancy for overall agent productivity. Additionally, while low agent occupancy guarantees that agent utilization will also be low, the reverse it not true: high agent occupancy does not guarantee a high agent utilization rate. Finally, when agent occupancy and utilization diverge – i.e., high occupancy and low utilization – this is generally an indication that the service desk is overstaffed.

Figure 1: Agent Occupancy vs. Agent Utilization

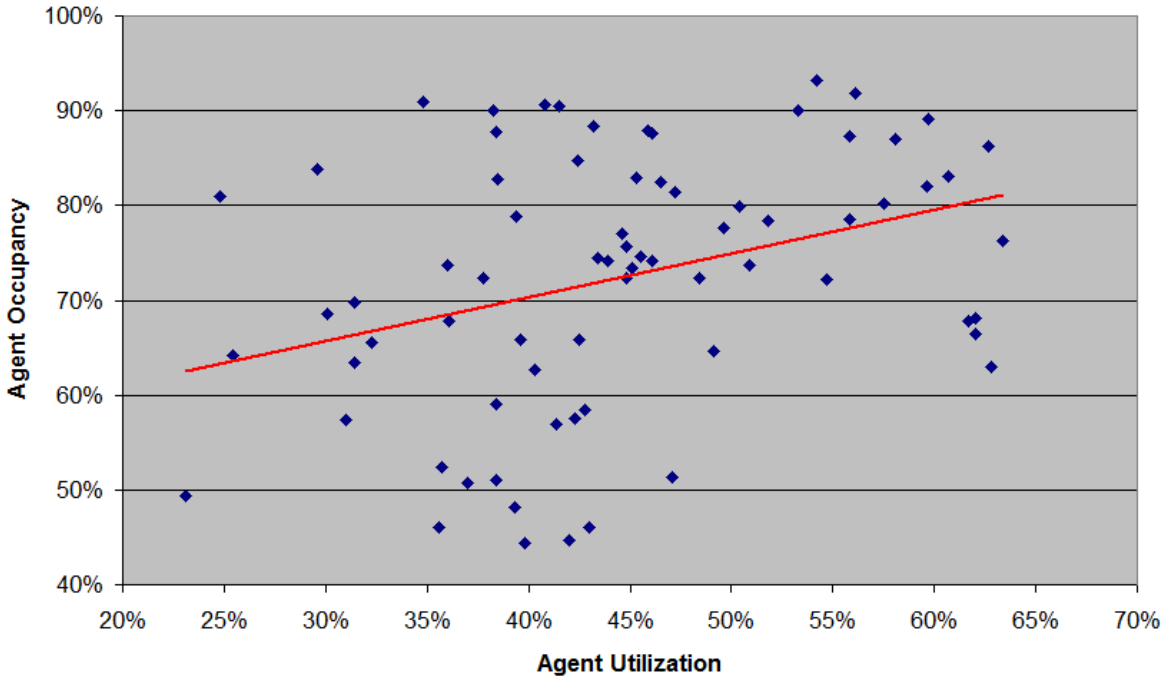

Benchmarking Ranges for Agent Occupancy

MetricNet's benchmarking database shows that the average agent occupancy for service desks worldwide is about 73%. This number varies widely, however, from a low of 44% to a high of 93%. Those at the low end of this range tend to be smaller service desks that are unable to achieve economies of scale, or service desks that are overstaffed. Those at the higher end of this range tend to be service desks that have good scale economies and good agent scheduling practices.

Figure 2 below shows the distribution of agent occupancy for a representative cross-section of service desks worldwide.

Figure 2: Benchmarking Statistics for Agent Occupancy

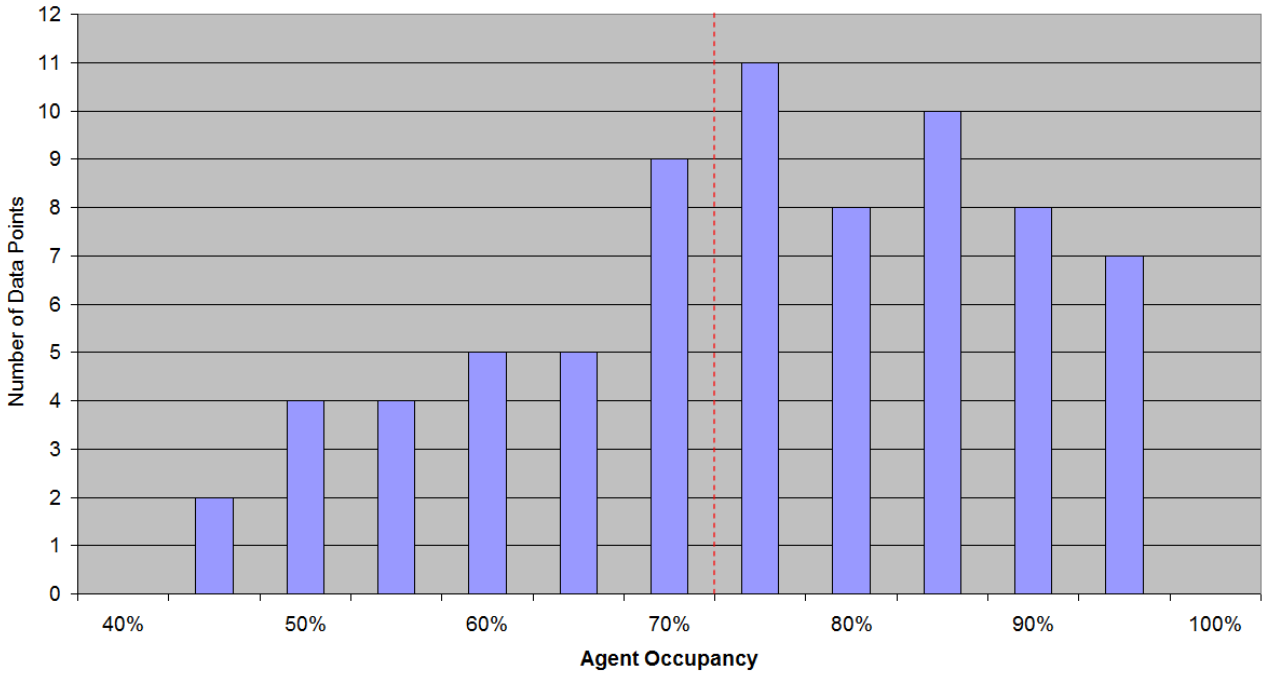

In summary, agent occupancy is a common and important service desk metric to track and trend. And it does provide an indication of how productive an agent is *while they are logged into the system*. But for the reasons outlined above, agent occupancy has limitations, and agent utilization is a better overall indicator of agent productivity.

Technician Utilization for Desktop Support

One goal of every business is to achieve the highest possible quality at the lowest possible cost. It stands to reason, therefore, that cost and quality should be measured on an ongoing basis. In fact, many would argue that cost and quality are the *only* two things that really matter in a service and support organization. In past articles MetricNet has discussed the importance of using metrics as a diagnostic tool to improve performance. So we have to ask ourselves, if cost per ticket is one of the foundation metrics for desktop support, how can we affect it? How can we improve it? What are the primary levers we have to manage cost?

Desktop support is a labor-intensive function. Technician salaries and benefits make up more than 60% of all costs for the average desktop support group. And when you consider the salaries and benefits for non-technicians – e.g. supervisors, team leads, QA/QC, trainers, and workforce schedulers – approximately 74% of all desktop support costs are personnel related, as shown in Figure 1 below. Labor productivity is therefore the biggest lever we have to manage and control desktop support costs.

Figure 1: Expense Breakdown for Desktop Support

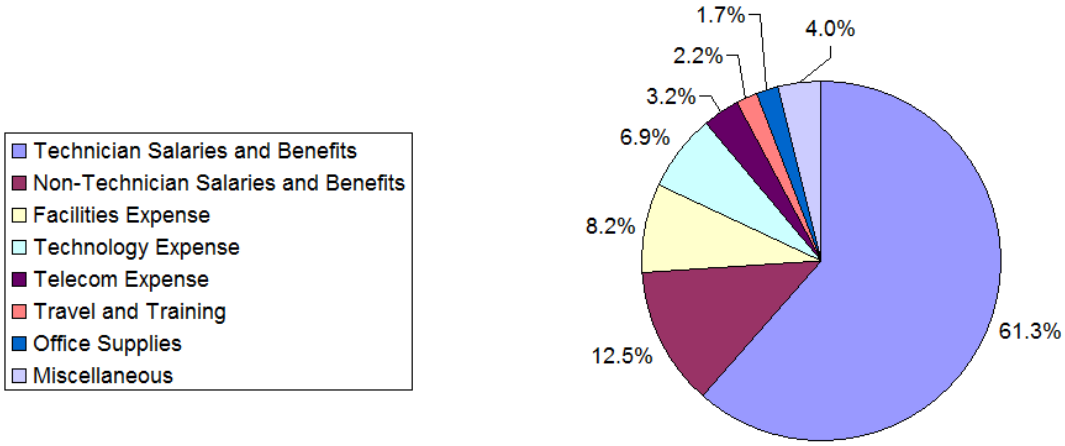

The best measure of labor productivity in desktop support is technician utilization. Because technician salaries and benefits represent more than 60% of all desktop support costs, if technician utilization is high, the cost per ticket will be correspondingly lower. Conversely, when technician utilization is low, technician costs, and hence cost per ticket, will be correspondingly higher. This is illustrated in Figure 2 below.

Figure 2: Technician Utilization vs. Cost per Ticket

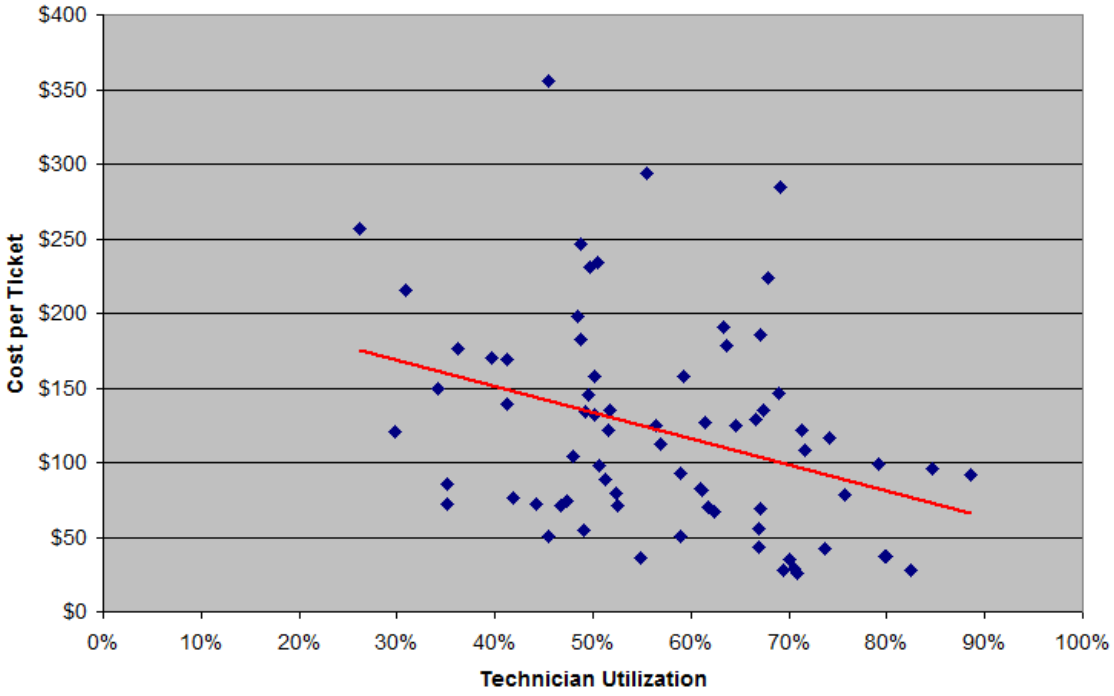

Just as world-class desktop support groups are obsessive about maintaining high customer satisfaction levels, they are equally committed to keeping their costs in check. They do this primarily by maintaining tight control over technician utilization. This has the effect of helping to minimize cost per ticket as illustrated above. Keep in mind, however, that other factors besides technician utilization will have an impact on the cost per ticket. Chief among these is the travel time per ticket, which can be quite high when supporting a customer base that is widely dispersed geographically.

The formula for determining technician utilization is somewhat complicated. It factors in the hours in a work day, break times, vacation and sick days, training time and a number of other factors. But there is an easy way to approximate technician utilization without going to so much trouble:

$$\text{Technician Utilization} = \frac{\begin{array}{l}((\text{Average number of Incidents handled by a technician in a month}) \times (\text{Average Incident Work Time}) + \\ (\text{Average number of Service Requests handled by a technician in a month}) \times (\text{Average Service Request Work Time}) + \\ (\text{Average number Tickets handled by a technician in a month}) \times (\text{Average Travel Time per Ticket}))\end{array}}{(\text{Average number of days worked in a month}) \times (\text{Number of work hours in a day}) \times (60 \text{ minutes/hr})}$$

Let's say, for example, that the technicians in a particular desktop support group handle an average of 60 incidents and 24 service requests per month. Let's further assume that the average handle time for incidents is 32 minutes, and the average handle time for service requests is 59 minutes. Let's also assume an average travel time per ticket of 41 minutes. Additionally, these technicians work an average of 21.5 days per month, and their work day is 7.5 hours after subtracting lunch and break times. The simplified utilization formula above would work out to the following:

$$\text{Technician Utilization} = \frac{\text{((60 Incidents per Month) X (32 minutes) + (24 Service Requests per Month) X (59 minutes) + (84 Tickets per Month) X (41 minutes))}}{\text{(21.5 working days per month) X (7.5 work hours per day) X (60 minutes/hr)}} = \text{70\% Technician Utilization}$$

Once again, this is not a perfect measure of technician utilization, but it is quick and easy, and gets you within 5% of the true technician utilization figure.

Benchmark Ranges for Technician Utilization

MetricNet's benchmarking database shows that the average technician utilization for desktop support groups worldwide is about 57%. This number varies widely, however, from a low of 26% to a high of 89%. Those at the low end of this range tend to be smaller organizations that are unable to achieve economies of scale, or desktop support groups that must overstaff to meet very aggressive service levels for response time. Those at the high end of this range tend to be larger organizations that have good scale economies, and can schedule their desktop support technicians very efficiently.

Figure 3 below shows the distribution of technician utilization for a representative cross-section of desktop support groups worldwide.

Figure 3: Benchmarking Statistics for Desktop Support Technician Utilization

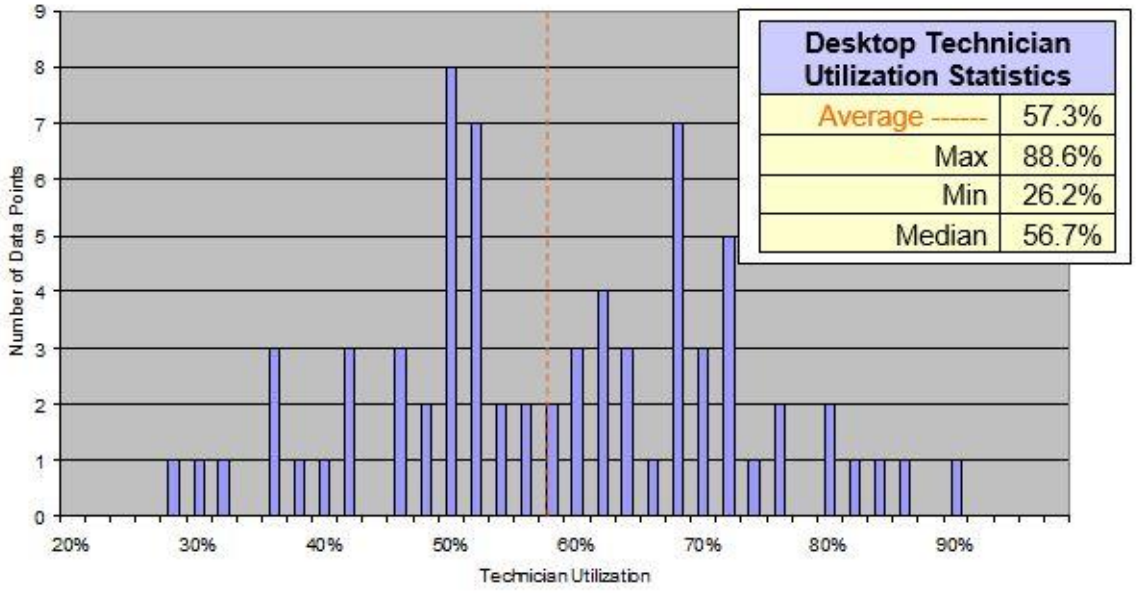

Desktop Support Balanced Scorecard

Today's technologies and reporting packages make it easy to capture copious amounts of performance data. Most support managers can tell you everything from last month's ticket volume to yesterday's mean time to resolve. But what does it all mean? If my cycle time goes up, but my cost per ticket goes down, is that good or bad? Is desktop support performing better this month than it was last month?

Despite all the data that support managers have at their fingertips, most cannot answer a very basic question: How is desktop support performing? The balanced scorecard resolves this dilemma by combining the most important desktop support KPI's into a single, overall measure of desktop support performance.

MetricNet's research shows that establishing an overall performance metric for desktop support is critical. We call this metric the balanced score because it truly does communicate a balanced picture of desktop support performance. The balanced scorecard is a mechanism that aggregates the most important desktop support metrics – such as cost per ticket and customer satisfaction – into a single, all-inclusive measure of desktop support performance. The value of this metric, when tracked over time, is that it enables desktop support to determine whether overall performance is improving or getting worse.

Oftentimes, when desktop support attempts to communicate its performance to other stakeholders in the business, particularly to lay people who do not understand desktop support operations, those people quickly become overwhelmed by the minutia of such measures as technician utilization, and they are confused about how to interpret the results. They are likely to focus in on one, easily-understood metric like customer satisfaction, and draw conclusions about overall desktop support performance from this single metric. This is a classic case of missing the forest for the trees. It is therefore absolutely critical to communicate the overall performance of desktop support, and the balanced scorecard does that for you. Think of the balanced scorecard as your letter grade for the month! In this way, desktop support can track its overall performance, and, in any given month, may see costs go up or customer satisfaction go down or mean time to resolve increase, but these individual measures take on a secondary level of importance because the balanced score provides a more complete and accurate picture of overall desktop support performance.

The Mechanics of Creating a Scorecard

Creating a scorecard is relatively straightforward. You can follow along in Figure 1 below as I explain the process.

Figure 1: Desktop Support Balanced Scorecard

Performance Metric	Metric Weighting	Performance Range		Co. XYZ Performance	Metric Score	Balanced Score
		Worst Case	Best Case			
Cost per Incident	15.0%	$288.04	$16.48	$61.00	83.6%	12.5%
Cost per Service Request	15.0%	$324.65	$35.60	$90.05	81.2%	12.2%
Customer Satisfaction	20.0%	59.1%	97.8%	84.8%	66.3%	13.3%
Incident First Visit Resolution Rate	10.0%	70.6%	97.1%	85.0%	54.3%	5.4%
Technician Utilization	10.0%	30.2%	72.6%	56.9%	63.0%	6.3%
% of Incidents Resolved within 8 work hours	10.0%	16.7%	95.1%	74.1%	73.2%	7.3%
% of Service Requests Fulfilled in 24 work hours	10.0%	7.0%	83.0%	52.2%	59.5%	5.9%
Technician Job Satisfaction	10.0%	67.8%	95.1%	79.4%	42.5%	4.3%
Total	**100.0%**	**N/A**	**N/A**	**N/A**	**N/A**	**67.2%**

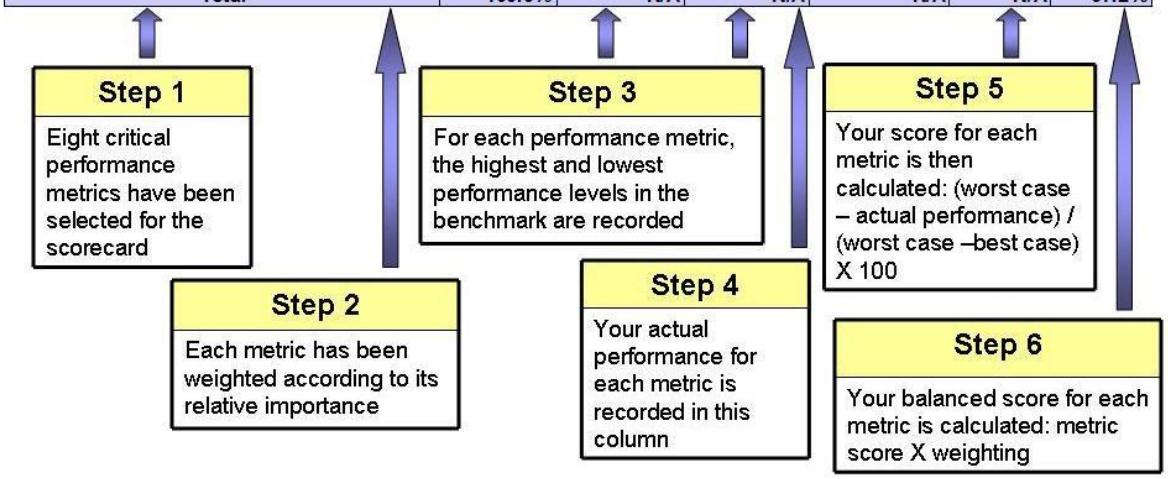

First you select the metrics to include in your scorecard. We suggest including the following eight metrics: Cost per incident, cost per service request, customer satisfaction, incident first visit resolution rate, technician utilization, % of incidents resolved within 8 work hours, % of service requests fulfilled in 24 work hours, and technician job satisfaction. Depending upon the metrics you track in desktop support, you may choose fewer metrics or a different mix of metrics for your scorecard. Secondly, you establish a weighting for each metric based upon its relative importance in the scorecard. This is a judgment call, but we suggest overweighting cost and customer satisfaction, since these are the foundation metrics for service and support. Step 3 is to show a reasonable range of performance – worst case to best case – for each metric. Normally these performance ranges are derived from a benchmark of desktop support. In step 4 your actual performance for each metric is inserted into the third column from the right. A score for each metric is then calculated based on the interpolation formula in step 5. And finally, a balanced score for each metric is determined by multiplying the metric weighting by the metric score. When the metric scores are summed up, you have the total balanced score for desktop support!

In this particular example, the desktop support balanced score is 67.2%. Your balanced score will always range from 0% (if you have the worst possible performance for every metric in the scorecard) to 100% (if you have the best possible performance for every metric in the scorecard). It turns out that the desktop support group in our example has scored quite well. When we run hundreds of desktop support groups through this algorithm, we get a normal distribution centered right at 50%. Those who score above 61% are in the top quartile; those who score between 50%

and 61 % are in the second quartile; those between 39% and 50% are in the third quartile; and those below 39% are in the bottom quartile for overall performance.

Benchmarking Your Performance

The balanced scorecard is an ideal way to track, trend, and benchmark your desktop support performance. Figure 2 below shows the trend in one company's desktop support performance over a 12 month period. The blue bars in the chart represent the monthly balanced scores, while the red background represents the 12 month trailing trend in scorecard performance. Clearly, the performance trend for this particular desktop support group is improving!

Figure 2: Balanced Scorecard Trend

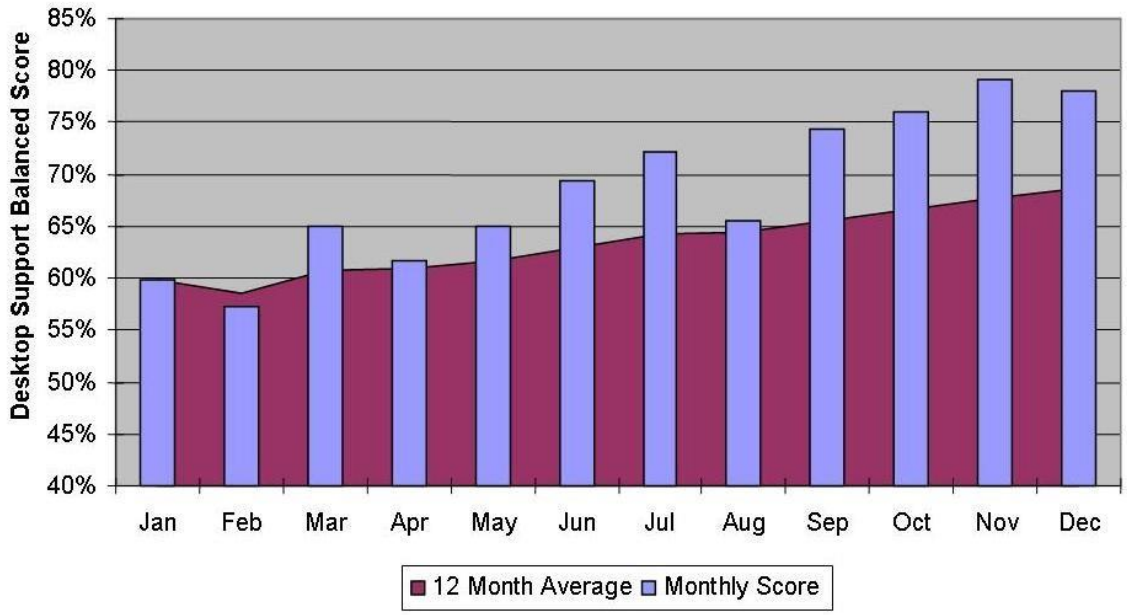

Finally, the desktop support balanced score can be used to benchmark your desktop support on a fair, apples-to-apples basis against other desktop support groups. Figure 3 below shows how the desktop support group in our example compares to 22 other desktop support groups in their benchmarking peer group.

Figure 3: Scorecard Benchmarking Comparison

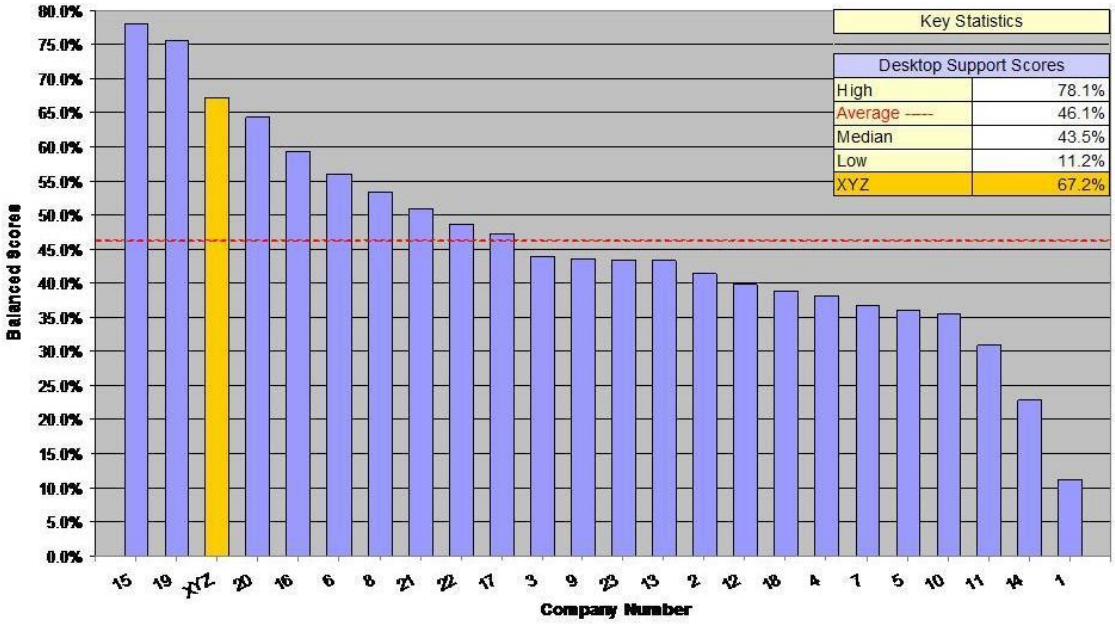

Self Service Completion Rate

The self service completion rate is a service desk metric that measures the percentage of level 1 incidents that are self-resolved by the user, without the assistance of a live agent. Let's say, for example, that the agents on a particular service desk handle 4,000 incidents per month through voice, chat, and email. Another 1,000 incidents per month are resolved through user self-service (e.g., through a password reset tool). For this hypothetical service desk the self-service completion rate is 1,000 self-service incidents ÷ 5,000 total incidents = 20% self-service completion rate.

The self-service completion rate is typically driven by the adoption of technologies that enable users to resolve their own incidents. The most common of these technologies include password reset tools, online FAQ's, and self-help portals that give users access to the solutions knowledge base. Some of these tools are quite powerful, and offer the users robust search capabilities for finding a solution to their specific need.

Why It's Important

Last year the average cost of resolving an incident at level 1 in North America was about $22. If some of these incidents can be deflected to self-service, that reduces the headcount requirements and hence the cost of the service desk. Years ago, some in the industry were predicting that user self-service tools would become so powerful that level 1 support would become obsolete. That, of course, did not happen, and it probably never will happen.

There are limitations on what can and should be resolved through self-service. As technology advances, incident complexity also increases. Developing knowledge articles that keep up with technology, and offering them in a friendly, searchable format for user self-service is a very challenging task. Moreover, you probably *don't* want your users spending a lot of time searching for solutions to their IT problems. Take the case of a user who spends two hours self-resolving a how-to issue for a proprietary business application. If that user could have gotten the same solution in 15 minutes by contacting the service desk, clearly that is the better support channel. This is the primary reason that most incidents resolved through self-service tend to be fairly straightforward – password resets, desktop application how-to's, and the like.

It is worth noting that as the self-service completion rate increases, organizations typically see an increase in the cost per agent-assisted contact. The reason for this is fairly obvious. The incidents resolved through self-service tend to have fairly short handle times when resolved by an agent. As these lower handle time incidents are deflected into the self-service channel, the average complexity, and hence the average handle time of the incidents that continue to be handled by live agents will increase. This, in turn, increases the average cost of an agent-assisted incident.

Benchmark Data for Self-Service Completion Rate

MetricNet's benchmarking database shows that the average self-service completion rate for service desks worldwide is 10.4%. This number varies widely, however, from a low of 0% for many service desks with no self-service capabilities, to a high of 55%. Not surprisingly, the vast majority of these self-resolved incidents are password resets. In fact, password reset tools are one of the most common technologies found in IT service and support, and have generally been shown to have a very favorable ROI.

Figure 1 below shows the distribution of self-service completion rate for a representative cross-section of service desks worldwide.

Figure 1: Self-Service Completion Rate

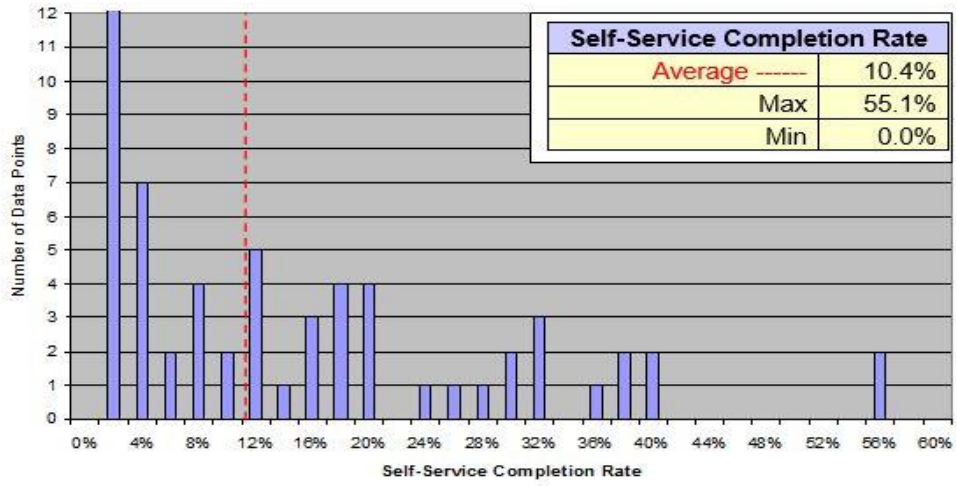

Chat Metrics | An Overview

This month we depart from our usual format, and instead of discussing a single metric, I will introduce and discuss a number of metrics that are unique to chat.

The Chat Channel

If you are like most consumers, you have probably experienced a chat session. Perhaps you engaged in chat with an agent at your bank or insurance company to resolve a payment issue. Or you may have used chat to troubleshoot your new computer or a software application you installed.

Why chat? One reason is that some people simply prefer this channel for service and support. The fact is, chat is the channel of choice for a growing number of consumers and businesses, particularly among millenial's. The second reason is economics. An effective chat channel can significantly reduce the cost per transaction vs. a more traditional live voice support model. Because of this, chat has the potential to both improve customer satisfaction (by giving customers an alternative channel choice), and to reduce cost per contact.

Costs are typically lower in the chat channel because chat agents generally handle more than one chat session at a time. These are called concurrent sessions, and some agents are talented and skilled enough to handle as many as four concurrent chat sessions. When agents handle more than one session at a time, the cost per session is reduced accordingly.

Chat Metrics

Many chat metrics simply mirror their live agent counterparts. For example customer satisfaction for chat is the same metric as customer satisfaction for the live voice channel. However, the *values* for the metrics in each channel are likely to be different. Customer satisfaction in the traditional voice channel might be 80%, while customer satisfaction for chat might be 90%, or vice versa.

The metrics listed below are common to both the chat and voice channels. Some names might vary slightly – for example, speed of answer in the voice channel becomes speed of response in the chat channel, and call abandonment rate becomes chat abandonment rate in the chat channel.

Cost Metrics
Cost per Contact
Cost per Minute of Handle Time

Productivity Metrics
Contacts per Agent per Month
Agent Utilization

Service Level Metrics
Average Speed of Answer / Average Speed of Response
Call/Chat Abandonment Rate
% Answered in 30 Seconds

Quality Metrics
Customer Satisfaction
First Contact Resolution Rate
Call/chat Quality

Agent Metrics
Agent Occupancy
Annual Agent Turnover
Daily Agent Absenteeism
New Agent Training Hours
Annual Agent Training Hours
Agent Tenure
Agent Job Satisfaction

Contact Handling Metrics
Inbound Contact Handle Time

There is another category of metrics that are unique to chat. These include the following:

% of Contacts Originating in Chat – as the name suggests, this is the percentage of all contacts coming into the service desk that originate in the chat channel. As a chat channel matures, this metric normally increases.

% of Contacts Resolved in Chat – this number will be less than or equal to the % of contacts originating in chat. Once again, as the chat channel matures, this metric normally increases.

% Failover Rate from Chat to Voice – this measures the percentage of chats that "failover" to a live agent voice contact. This happens when the agent or caller feels that voice communication is needed, and they revert from the chat channel to the voice channel to complete a transaction.

Average Concurrent Chat Sessions – this metric indicates the average number of concurrent sessions a chat agent has open at any given time. The ability to handle concurrent chat sessions is the primary economic advantage of the chat channel.

Maximum Concurrent Chat Sessions – Most organizations will limit the number of concurrent sessions an agent is allowed to handle. Newer agents might be limited to a single chat session at a time, while more experienced agents might be allowed to handle as many as four concurrent chat sessions.

Service desks that are interested in growing, improving, and optimizing chat can use these metrics to effectively manage the chat channel.

Benchmarking Data for Chat

Chat data from MetricNet's global benchmarking database is excerpted below in Figure 1. Please note that only averages are shown; the range for each metric – from high to low – can be quite large and dramatic.

Figure 1: Benchmark Averages for Chat Metrics

Metric Type	Chat KPI's	North American Averages
Cost	Cost per Inbound Chat Contact	$11.20
	Cost per Minute of Inbound Chat Handle Time	$0.90
Productivity	Inbound Chats per Technician per Month	670
	Chat Technician Utilization	54.1%
	Chat Technicians as a % of Total Headcount	84.3%
Service Level	Average Chat Speed of Answer (ASA) (minutes)	3.93
	Chat Call Abandonment Rate	13.1%
Quality	Chat Call Quality	80.2%
	Chat First Contact Resolution Rate	70.2%
	Chat Customer Satisfaction	79.5%
	Chat Net Promoter Score	63.1%
Technician	Annual Chat Technician Turnover	22.9%
	Daily Chat Technician Absenteeism	8.0%
	Chat Technician Occupancy	73.0%
	New Chat Technician Training Hours	162
	Annual Chat Technician Training Hours	30
	Chat Technician Tenure (months)	52.2
	Chat Technician Job Satisfaction	75.8%
Chat Handling	Inbound Chat Handle Time (minutes)	12.90
	% of Contacts Originating in Chat	10.9%
	% of Contacts Resolved in Chat	6.6%
	% Failover Rate from Chat to Voice	32.0%
	Average Concurrent Chat Sessions	0.84
	Max Concurrent Chat Sessions	2.75

Ideally the chat channel should enrich the user experience by providing channel choice, and high quality transactions. However, it should be noted that the agent skill set required for chat is somewhat different than that required of a live voice agent. You should not automatically assume that a successful voice agent will be a successful chat agent, and vice versa.

Ratio of Agents to Total Headcount

There are a number of ways to measure the efficiency of a service desk or desktop support group. Metrics such as cost per ticket and agent utilization are the most common measures of efficiency. But there is another, less well known metric that also drives cost per ticket. That metric is the ratio of agents to total headcount, and it applies equally to both service desk and desktop support groups.

By way of example, let's assume that a service desk has 25 front-line agents, 3 supervisors, 1 QA/QC person, and 1 person for training and workforce scheduling. That's a total headcount of 30. In this example, the ratio of agents to total headcount is 25 ÷ 30 = 83.3%. The same calculation would apply for a desktop support group.

The ratio of agents to total headcount is an indirect measure of managerial efficiency in service and support. A high ratio of agents to total headcount is indicative of lean management staffing, while a low ratio of agents to total headcount is symptomatic of a top-heavy organization. It is, however, possible for the ratio to be too high, in which case the support organization may lack adequate headcount for proper supervision and other indirect functions such as reporting, workforce scheduling, and training. Conversely, there are times when a low ratio of agents to total headcount may be justified, at least temporarily, such as when a support group moves agents out of customer facing roles to work on and contribute to various projects.

Benchmarking Data for Ratio of Agents to Total Headcount

When a support organization experiences costs that are higher than the industry average, it can sometimes be attributed to the ratio of agents to total headcount. In general, a high ratio (meaning low indirect headcount) will lead to lower costs, while a low ratio (meaning high indirect headcount) will lead to higher costs. Indeed, data for this metric extracted from MetricNet's benchmarking database and illustrated in Figure 1 below, shows a downward slope in Cost per Ticket for both the service desk and desktop support as the ratio of agents to total headcount increases.

Figure 1: Ratio of Agents to Total Headcount vs. Cost per Ticket

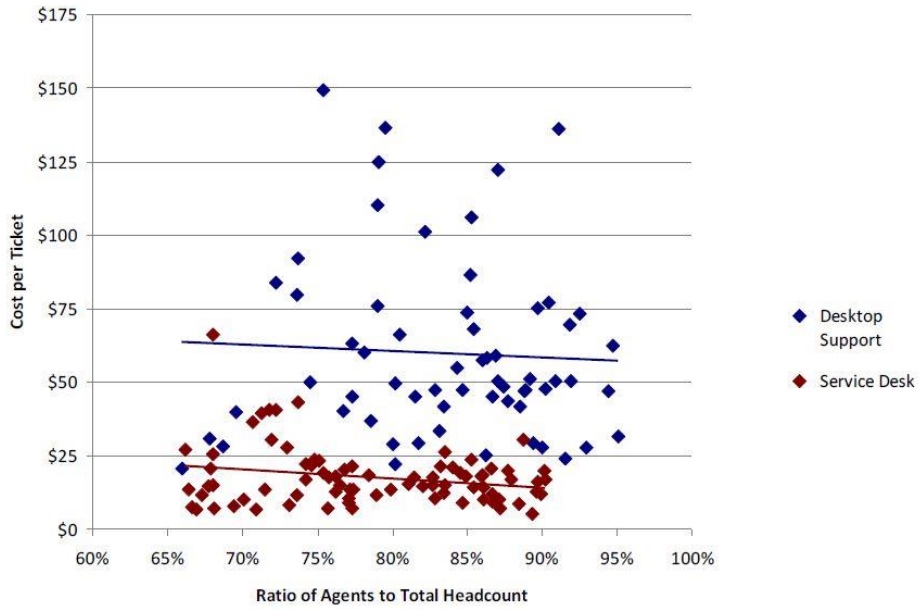

Figure 2 below shows the distribution of agents to total headcount for both the service desk and desktop support, along with summary statistics for both metrics. The ratio of agents to total headcount can vary quite widely, from 66% to 90% for the service desk, and from 66% to 95% for desktop support. The average ratio of agents to total headcount tends to be a bit higher for desktop support than for the service desk. This reflects the fact that supervision, reporting, training and scheduling in desktop support is oftentimes less rigorous than in the service desk.

Figure 2: Distribution of Agents to Total Headcount

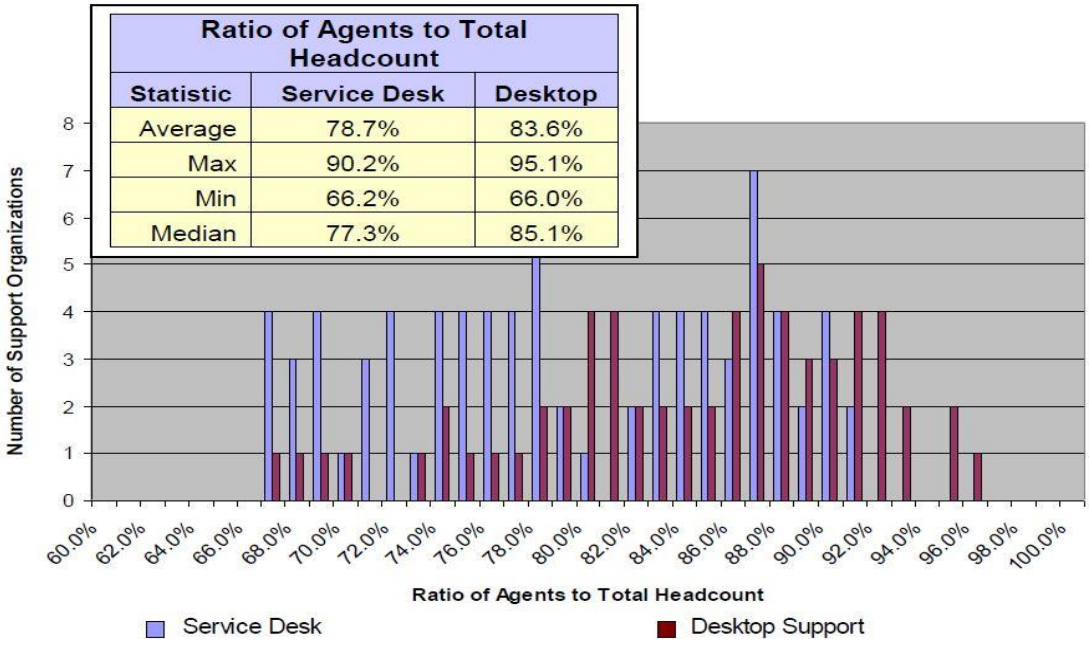

ROI of Support | Part I

This article begins a two part series on Return on Investment for service and support. In part 1 I define how value is created in IT service and support. In part 2, I will go through a case study that calculates the ROI for a particular support organization.

ROI of Support

Most IT departments can tell you how much they spend on support. But very few can quantify the economic impact of support. The result is that many IT service and support organizations are on the defensive when it comes to budgeting and spending, and often struggle just to get the funding needed to deliver adequate levels of support.

In recent years a handful of pioneering organizations have adopted a different strategy when it comes to support – a strategy that emphasizes value over cost – and they routinely deliver benefits far in excess of their costs. Support groups that understand and quantify their ROI gain a number of important advantages; chief among them is the ability to obtain funding and other resources based upon the economic benefits of the support they deliver.

IT support has the potential to deliver positive ROI in at least two ways. The first source of value derives from the ability of support to minimize total cost of ownership (TCO). Last year the average cost per ticket was $22 for the service desk, $62 for desktop support, and $85 for level 3 IT support as shown in Figure 1 below. The well-known industry expression "Get it done at Level 1" is not just a pithy catchphrase…it can literally save your organization millions of dollars a year!

Recent benchmarks show that 21% of all tickets resolved by desktop support could have, and should have been resolved by the level 1 service desk. These escalation errors represent defects in the support process, and result in increased support costs that often go unnoticed because they are rarely tracked. By investing in tools, training, and technology that drive higher resolution rates at level 1, you have the potential to significantly reduce the TCO for IT support.

Figure 1: Cost per Ticket – North American Averages

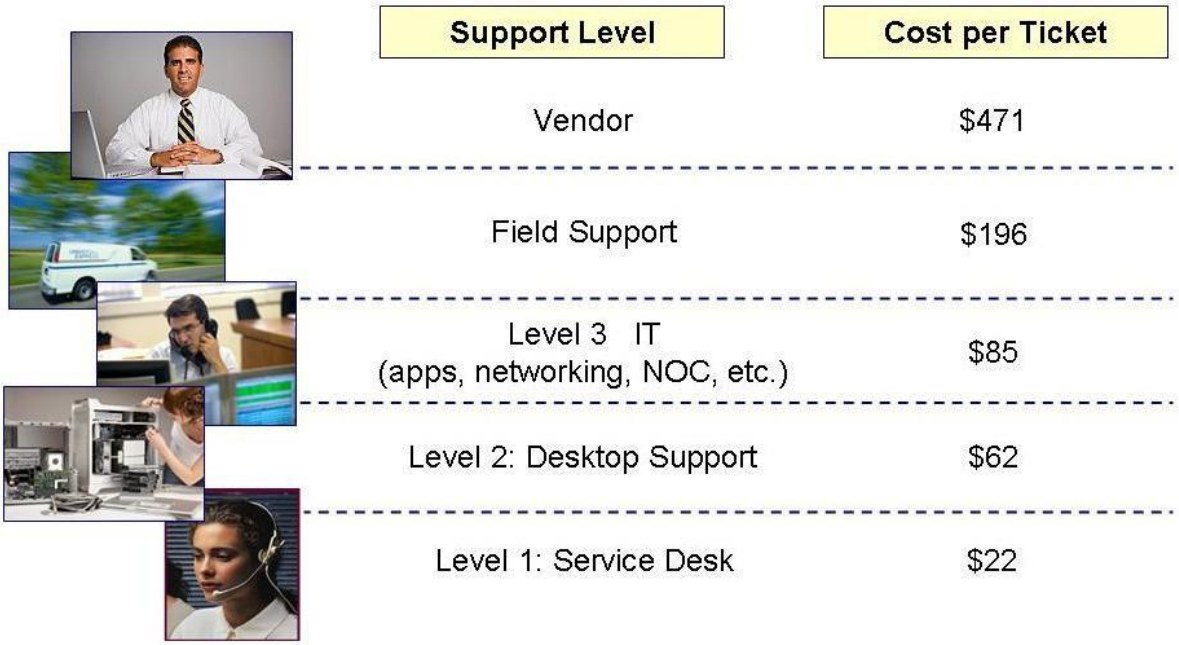

The second source of value creation and ROI is the result of returning productive hours to the workforce. The majority of today's workforce is comprised of knowledge workers, all of whom rely upon one or more computing devices to do their jobs. When these devices break down or do not function properly, employee productivity suffers. By preventing these incidents from occurring, and by quickly resolving issues when they do occur, a support organization can return productive hours to the workforce.

A study conducted by MetricNet and summarized in Figure 2 below concluded that knowledge workers lose an average of 33 hours of productive time per year due to various IT outages, breakdowns, and hardware and software failures. For support groups performing in the top quartile, the lost productivity per worker is just 17 hours per year; about half the industry average. By contrast, employees who receive support from bottom quartile support groups lose an average of 47 productive hours per year.

The difference between the top and bottom quartile performers is a staggering 30 hours per employee per year! Put another way, support organizations in the top quartile are able to return nearly four extra days of productivity annually for every knowledge worker in the enterprise. When multiplied by thousands or even tens of thousands of employees in a company, the productivity gains and ROI delivered by a top performing support group can be enormous!

Figure 2: Economic Benefits of Enhanced User Productivity

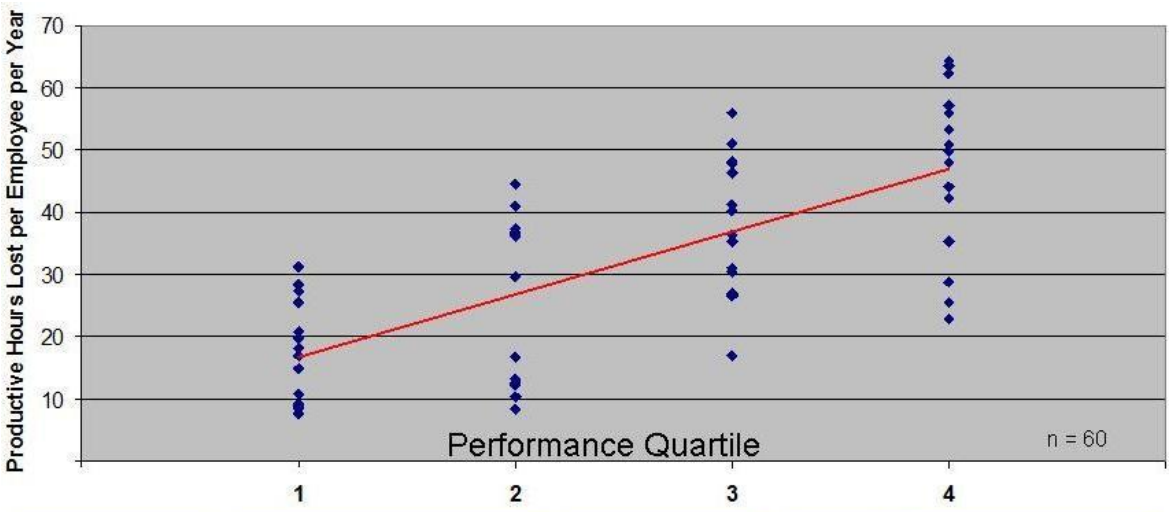

Support Function	Key Performance Indicator	Performance Quartile			
		1 (top)	2	3	4 (bottom)
Service Desk	Customer Satisfaction	93.5%	84.5%	76.1%	69.3%
	First Contact Resolution Rate	90.1%	83.0%	72.7%	66.4%
	Mean Time to Resolve (hours)	0.8	1.2	3.6	5.0
Desktop Support	Customer Satisfaction	94.4%	89.2%	79.0%	71.7%
	First Contact Resolution Rate	89.3%	85.6%	80.9%	74.5%
	Mean Time to Resolve (hours)	2.9	4.8	9.4	12.3
Average Productive Hours Lost per Employee per Year		17.1	25.9	37.4	46.9

Now that you understand how value (the return in ROI) is created in IT service and support, the next step is to apply this knowledge to calculate the ROI of your support organization.

What We Do

MetricNet is the leading source of benchmarks, scorecards, and performance metrics for Information Technology and Call Center Professionals worldwide. Our mission is to provide you with the benchmarks you need to run your business more effectively. MetricNet has pioneered a number of innovative techniques to ensure that you receive fast, accurate benchmarks, with a minimum of time and effort.

We offer **The One Year Path to Excellence**, a continuous improvement program, **industry benchmarks** that are downloadable from our website, **benchmarking data files** for those who wish to conduct their own benchmarking analysis, and comprehensive **peer group benchmarks** that compare your performance to others in your vertical market.

Our Clients ▸
and benchmarking partners

SONY · American Express · Progress Energy · PURINA · Desjardins · Symantec · WELLSTAR · GM · Schneider Electric · Bayer · Allstate (You're in good hands) · ARAMARK