The IT Service & Technical Support Community

Staffing Desktop Support

How Many Technicians Do You Need?

by

Jeff Rumburg
Managing Partner | MetricNet, LLC

UBM TechWeb

Contents

Introduction . . . 3

The Staffing Fallacy in Desktop Support . . . 3

Tickets, Incidents, and Service Requests . . . 5

Desktop Staffing Fundamentals . . . 6

What Is Technician Utilization? . . . 7

Metrics Needed for Staffing Decisions . . . 12

Optimizing Technician Headcount . . . 13

Conclusion . . . 14

Appendix: What Is a SPOC Service Desk? . . . 15

About the Author . . . 17

About MetricNet, LLC . . . 17

About HDI . . . 17

Introduction

One of the most common questions I hear from IT support managers is "How many technicians should I have in desktop support?" It's a great question, but one that is rarely answered adequately. The result is that many desktop support organizations are not staffed properly because they do not follow any sort of proven methodology for making headcount decisions. Instead, they rely on "gut feel" or "instinct" when it comes to staffing desktop support. Compounding this situation is the fact that many desktop support groups do not follow a strict SPOC (single point of contact) support model, and end up handling large numbers of incidents that could and should be resolved by the service desk. It's no wonder then that many desktop support groups are overstaffed and, consequently, very costly.

In this white paper, I will present a rigorous methodology for determining the appropriate technician headcount for desktop support. By following this approach, desktop support organizations can be assured that they will be staffed to meet the needs and expectations of their customers, while simultaneously delivering services in an efficient, fiscally responsible manner.

The Staffing Fallacy in Desktop Support

A common misconception in desktop support is that the user population alone will define the number of technicians needed. This approach wrongly assumes a fixed ratio of desktop support technicians to the number of users supported; for example, 12.5 desktop support technicians for every 1,000 users. The error in this approach is that no two user populations have the same needs; therefore, no two user populations generate the same workload. As such, staffing decisions in desktop support should be based upon workload, not user population.

Consider the example of a financial services company with a corporate staff of 2,500. The device count for this user population is as follows:

- 2,100 desktop computers
- 950 laptop computers
- 140 printers and copy machines
- 1,100 smartphones
- 240 servers

The number of desktop support technicians required to support this workload is about twenty-four. (The derivation of this number will be shown later in this paper.)

Now take the example of an electric utility company with the same number of users (2,500). The device count for this population is as follows:

- 2,200 desktop computers
- 550 laptop computers
- 80 printers and copy machines
- 300 smartphones
- 70 servers

The number of desktop support technicians required to support this workload is about eight. (Again, the derivation of this number will be shown later in this paper.)

Despite the fact that the user populations in our examples are exactly the same, the two user groups generate very different workloads and therefore require different desktop support staffing levels. The workload for desktop support is driven not only by the number of users or "seats supported," but also by the number and mix of devices, the age of the devices, the organization's service level targets, the user population density, the level of standardization and virtualization in the IT environment, and a myriad of other factors.

The key point here is this: The technician staffing requirements of a desktop support organization should be based upon the workload generated by the user population—incident and service request volume—not the number of users being supported. With this in mind, it is easy to understand why two organizations with exactly the same headcount may require very different desktop support staffing levels.

Tickets, Incidents, and Service Requests

Before we go any further, it is worth defining a few terms. Specifically, we need to define the terms *ticket*, *incident*, and ***service request*** as they relate to desktop support.

An **incident** is typically unplanned work that requires the assistance of an on-site technician to resolve. Some 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 (L1) service desk. By contrast, most **service requests** represent planned work. Some of the most common service requests are moves/adds/changes, hardware refresh/replacement, and device upgrades. **Tickets** represent the sum of all incidents and service requests. Figure 1 illustrates these definitions.

Figure 1: Tickets, Incidents, and Service Requests

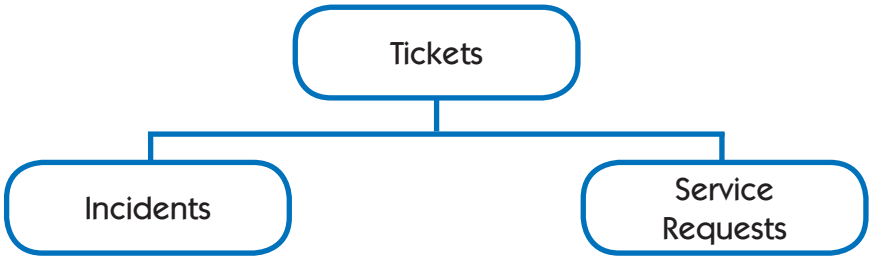

Unplanned work that requires a physical touch to a device:

- Hardware break/fix
- Device failure
- Connectivity failure

Planned work that requires a physical touch one or more devices:

- Move/add/change
- Hardware or software upgrade
- Device refresh or set up

Incident Volume + Service Request Volume = Ticket Volume

Desktop Staffing Fundamentals

Since staffing is based on workload, it makes sense to first define what we mean by workload. Workload is the sum of all incidents and service requests multiplied by their respective handle times:

$$\text{Monthly Workload} = \left[\text{Incident volume per month} \times \text{Incident handle time}\right] + \left[\text{Service request volume per month} \times \text{Service request handle time}\right]$$

Handle time includes both work time and travel time, where travel time is the average round-trip time it takes the technician to get to and from the site of an incident or service request (more on this later).

Let's perform the workload calculation for the financial services example from above. The 2,500 employees at the corporate headquarters generate an average of 1,750 incidents and 800 service requests per month. The average travel time to and from the site of an incident or service request is about 25 minutes. This is considered a medium-density user environment, consistent with a corporate campus containing multiple office buildings. Additionally, the average work time for an incident is 18 minutes, while the average work time for a service request is 51 minutes.

We can now calculate the incident and service request handle times. The incident handle time is 43 minutes (25 minutes travel time + 18 minutes work time). The service request handle time is 76 minutes (25 minutes travel time + 51 minutes work time). With that additional data, we can calculate the workload using the formula above:

$$\left[1{,}750 \text{ incidents per month} \times 43 \text{ minutes handle time per incident}\right] + \left[800 \text{ service requests per month} \times 76 \text{ minutes handle time per service request}\right] = 136{,}050 \text{ workload minutes per month}$$

Now that we have the workload for our financial services company, if we know how many minutes the average technician works in a month, we can determine the technician headcount. Based on data from MetricNet's worldwide Desktop Support Benchmarking Database, the average desktop support technician logs about 5,800 work minutes per month (that is, the sum of all the incidents and service requests handled by a desktop support technician in a month, multiplied by the incident and service request handle time). This number is called *technician capacity*, and it equates to a technician utilization rate of 60 percent, the approximate industry average for technician utilization. This figure is the net of all vacation days, sick days, training time, administrative time, and other "nonworking" time.

What Is Technician Utilization?

Desktop support is a labor-intensive function. For the average desktop support organization, 74 percent of all costs are labor-related: salaries, benefits, incentives, and contractors. The best measure of labor efficiency is *technician utilization*. When technician utilization is high, the cost per ticket will be low. Conversely, when technician utilization is low, labor costs, and thus cost per ticket, will be high. But what is technician utilization, and how is it calculated?

Technician utilization is the average time a technician spends handling both incidents and service requests per month, divided by the number of work hours in a given month. It measures the percentage of time that the average technician is in "work mode," and is independent of ticket work time or complexity. The formula for technician utilization and a sample calculation are shown in Figures 2 and 3 below.

Figure 2: Technician Utilization Defined

[Average number of incidents handled by a technician per month × Average incident work time

+

Average number of service requests handled by a technician per month × Average service request work time

+

Average number of tickets handled by a technician per month × Average travel time per ticket] ÷ [Average number of days worked per month × Number of work hours per day × 60 minutes/hour]

Definition-at-a-Glance:

- It is a measure of technician work and travel time, divided by total time at work during the month.
- It takes into account both the incidents and service requests handled by technicians.
- It does not adjust for sick days, holidays, training time, project time, or idle time.

Figure 3: Sample Desktop Support Technician Utilization Defined

Parameters:

- Incidents per technician per month = 60
- Service requests per technician per month = 24
- Average tickets per technician per month = 84
- Average incident work time = 32 minutes
- Average service request work time = 59 minutes
- Average travel time per ticket = 41 minutes

[60 incidents per month × 32 minutes + 24 service requests per month × 59 minutes + 84 tickets per month × 41 minutes] ÷ [21.5 days worked per month × 7.5 work hours per day × 60 minutes/hour] = 70%

We can now calculate our desktop support technician headcount requirements as follows:

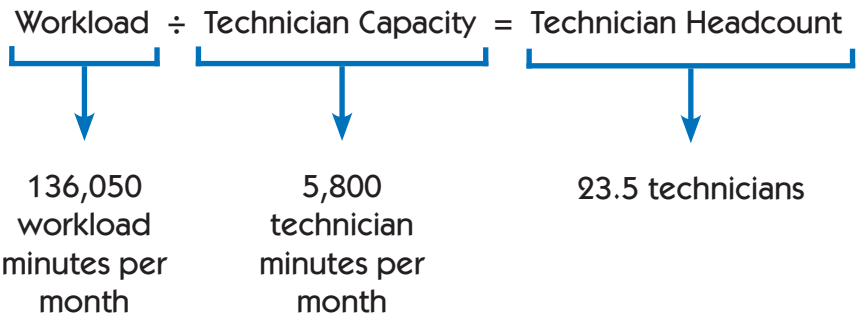

To properly staff the financial services desktop support function, we would need about twenty-four technicians.

How does this methodology work out for the electric utility example? The 2,500 employees at the utility headquarters generate an average of 900 incidents and 500 service requests per month. The average travel time to and from the site of an incident or service request is about 12 minutes. This is considered a high-density user environment, consistent with a high-rise office building. Additionally, the average work time for an incident is 14 minutes, while the average work time for a service request is 33 minutes. Thus, the incident handle time is 26 minutes (12 minutes travel time + 14 minutes work time), and the service request handle time is 45 minutes (12 minutes travel time + 33 minutes work time).

We can now calculate workload using the formula from p. 6:

[900 incidents per month × 26 minutes handle time per incident] + [500 service requests per month × 45 minutes handle time per service request] = 45,900 workload minutes per month

Finally, using a technician capacity of 5,800 work minutes per month, we can calculate our headcount requirements as follows:

45,900 workload minutes per month ÷ 5,800 technician minutes per month = 7.9 technicians

So we would need about eight technicians to properly staff the desktop support function at the electric utility.

These two examples clearly illustrate the importance of basing headcount decisions on workload, rather than the size of the customer population. In the case of the financial services company, we have twenty-four technicians supporting 2,500 users. This is one technician for every 104 users. By contrast, we have eight technicians supporting 2,500 users at the electric utility. This is one technician for every 313 users.

The table below, extracted from MetricNet's global Desktop Support Benchmarking Database, illustrates just how dramatically the workload, and hence headcount requirements, can vary from industry to industry, and even within a particular industry. Figure 4 (on the following page) illustrates those ranges graphically.

KPI	Statistic	Financial Services	High Tech	Equipment Manufacturing	Telecom	Business Services	Healthcare	Energy Utilities
Incident Work Time	Avg.	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	Avg.	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	Avg.	25	32	33	19	27	20	12
	Min.	11	17	14	9	11	8	7
	Max.	110	153	164	64	79	52	53
Incidents per Seat per Month	Avg.	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	Avg.	0.32	0.41	0.21	0.31	0.42	0.10	0.20
	Min.	0.10	0.29	0.13	0.13	0.13	0.04	0.10
	Max.	1.19	1.41	0.62	0.93	1.43	0.36	0.59
Desktop Techs per 1,000 Seats Supported	Avg.	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 4: Desktop Support Technician Headcount Ranges by Industry

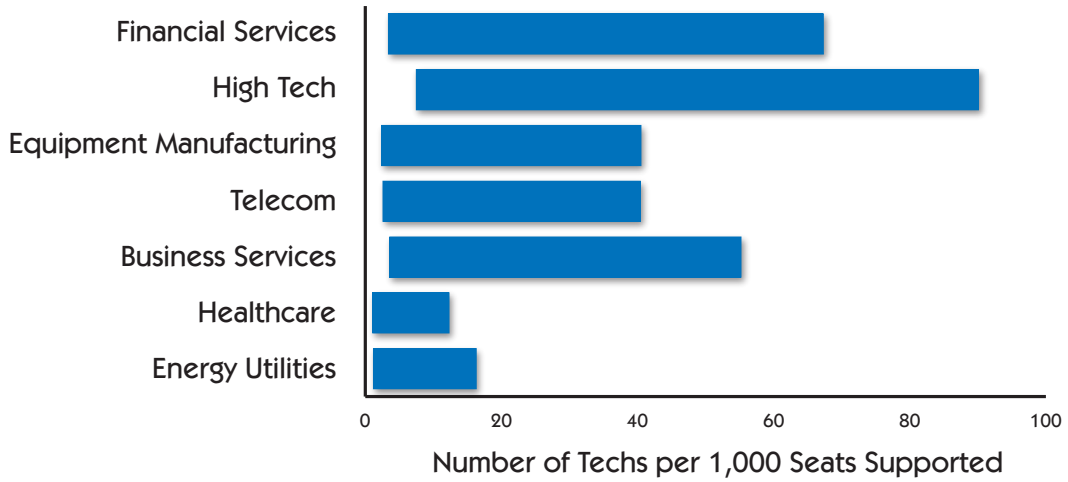

The variation in headcount requirements is driven by the wide range of workload requirements from company to company and from industry to industry. This, in turn, is a function of travel time, and the number of tickets generated by the user population. Travel time can range from just a few minutes in some high-density user environments (think high-rise office building), to as much as two hours or more in a low-density user environment (multiple buildings spread out over several hundred square miles, for example). The number of tickets generated, as mentioned before, is driven by numerous factors, including the mix of laptop and desktop computers, the number of remote users, the number of mobile devices, the device refresh rate, the standardization (or lack thereof) of the IT environment, and the degree of virtualization. Figures 5 and 6 below show just how dramatically incident and service request volume can vary.

Figure 5: Monthly Incident Volume by Industry

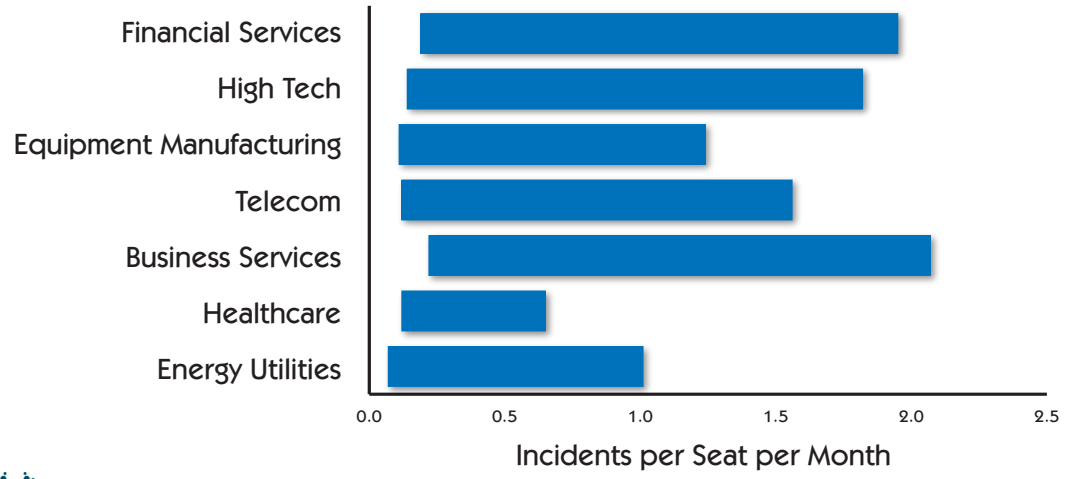

Figure 6: Monthly Service Request Volume by Industry

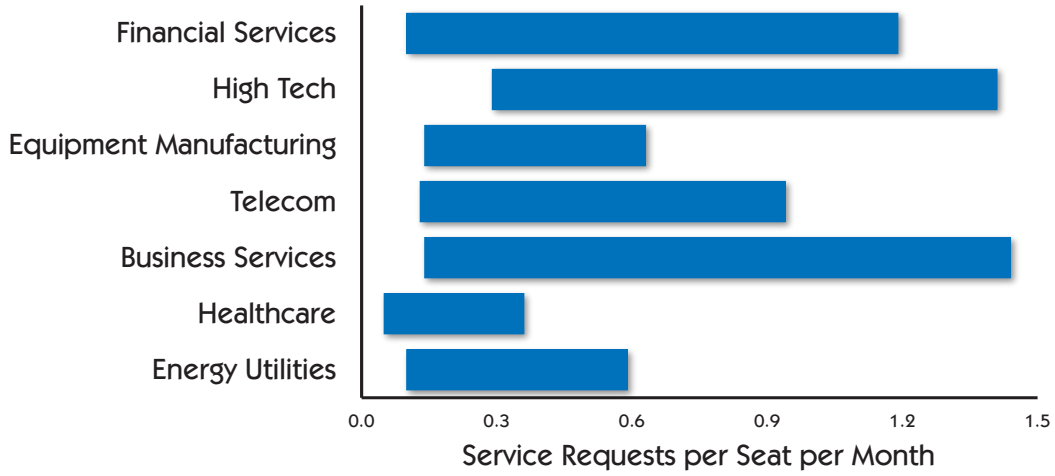

Metrics Needed for Staffing Decisions

In the July/August 2010 *SupportWorld* article entitled "Best Practices in Desktop Support: The Eight Essential KPIs for World-Class Performance," MetricNet noted that most desktop support organizations do not have a well developed set of metrics for tracking and trending their performance. However, the headcount methodology outlined above relies upon the following metrics:

- Monthly incident volume
- Monthly service request volume
- Average work time per incident
- Average work time per service request
- Average travel time per ticket

While most desktop support organizations track incident and service request volume, fewer than 20 percent track their work time or travel time. This presents a dilemma, but it's one that can be easily overcome. The simplest and most straightforward solution I have seen is when an entry for work time and travel time are added to the trouble ticket. When a technician closes out a ticket, he or she will estimate the work and travel time for that particular ticket, and enter those values into the appropriate fields on the ticket. The work and travel time for incidents and service requests can then be totaled up and averaged each month to develop a pretty good estimate for these numbers. These metrics, in turn, feed the calculations for determining technician headcount.

Some organizations have been known to engage in time and motion studies whereby the technicians in desktop support will carry stopwatches with them for a certain period of time (say, a week). During this time, each technician uses the stopwatch to record the precise work time and travel time for every ticket they handle. At the end of the measurement period, incident work time, service request work time, and travel time are totaled up and averaged for all technicians in desktop support. These values are then used to make or validate staffing decisions, calculate technician utilization, and update the desktop support balanced scorecard. For a detailed treatment of technician utilization and the balanced scorecard, please see the aforementioned *SupportWorld* article on desktop support KPIs. Finally, the time and motion study should be repeated once or twice a year to ensure that the values for work and travel time keep pace with the changing IT environment.

Finally, larger and more sophisticated desktop support organizations will sometimes equip their technicians with handheld devices that can be placed in various modes, such as work, travel, administrative, training, standby, phone support, etc., that keep track of individual work and travel time for both incidents and service requests. This, of course, makes the task of tracking these metrics much simpler and more accurate.

Optimizing Technician Headcount

Our discussion thus far has focused on a methodology that bases technician headcount on desktop support workload. But there is a second critical factor that must be taken into account when staffing a desktop support organization. Specifically, when a desktop support organization does not follow a strict SPOC support model, they often handle a large number of incidents that should be resolved by the L1 service desk. (For more information on the SPOC support model, see the appendix.) This creates inefficiencies in the support model and drives up the cost of support.

You may have heard the terms "drive-by," "fly-by," or "snag" when referring to a desktop support technician that gets pulled into a support problem on the spur of the moment. This might happen, for example, when a technician is returning to her desk after completing a service request, and is "snagged" by a user who needs support for a particular issue, right then and there. These rogue requests happen in every organization, and there is a great temptation for technicians to provide support when asked, even if it means violating the SPOC support protocol. This effect, sometimes called *bypass* (because it bypasses the SPOC process), can actually be measured with a metric called *percent resolved level 1 capable*, which, as the name implies, measures the number of tickets resolved by desktop support that could have been resolved by the service desk.

Globally, this metric averages 21 percent. That is, 21 percent of all tickets resolved by desktop support had the potential to be resolved by the L1 service desk. Some of these

tickets are unnecessarily escalated by L1 to desktop support, while others are the result of bypass, as described above. Either way, they represent defects in the support delivery model, and significantly inflate the cost of support.

In North America last year, the average fully-burdened cost of resolving a ticket at L1 was about $21, while the average fully-burdened cost of resolving a ticket at desktop support was about $63. For a desktop support organization handing 2,500 tickets per month, 21 percent L1 capable equates to 525 tickets that should have been resolved at L1. The average cost of these defects is $22,050 (or, 525 tickets × [$63 per desktop ticket – $21 per service desk ticket]). That's $22,050 in unnecessary support costs each month, or $264,600 per year!

The implication here is that the measured workload for desktop support can be inflated by tickets that should be resolved at L1. When this happens, the desktop support organization may appear to be staffed properly, when in fact it is overstaffed because it's handling a large number of escalated and bypass tickets that should have been resolved at L1. As shown above, this L1 "subsidy" can be quite costly. The solution to this problem is to track this metric – percent resolved level 1 capable – and make every effort to minimize ticket defects by enforcing a strict SPOC support model, and by providing the training and tools necessary at L1 to minimize the number of tickets that are unnecessarily escalated to desktop support.

Conclusion

Most desktop support organizations lack a rigorous methodology for determining technician staffing levels. Additionally, many do not follow a strict SPOC support model. These factors often result in staffing levels that are far from optimal, and can saddle a support organization with high costs.

In this article, we have presented a straightforward methodology for determining the appropriate technician headcount for desktop support. Additionally, we have recommended tracking a handful of metrics, including percent resolved level 1 capable, that can help to ensure that your desktop support function is staffed efficiently. Finally, we have shown that following a strict SPOC support model is crucial to achieving optimal staffing levels in desktop support.

Appendix: What Is a SPOC Service Desk?

The term SPOC refers to a single point of contact support organization. This means that all IT issues, service requests, problems, and incidents are first directed to the L1 service desk to be logged in the ticketing system, and then either resolved at L1 or dispatched to another source of support where the ticket can be resolved.

A SPOC service desk is not expected to resolve every ticket it logs. Rather, a SPOC service desk is a facilitator and coordinator of the entire end-user support process. They are responsible for resolving the tickets that can be resolved at L1, and expeditiously dispatching tickets that cannot be resolved at L1 to the most appropriate source of support. This could be desktop support, L2 IT groups, the NOC, a vendor, or even specific individuals in the organization with unique expertise, say, for a particular application. Finally, a SPOC monitors the progress of all open tickets, prompting action on tickets that appear to be stalled, and closing tickets that have been resolved satisfactorily. A SPOC service desk is ultimately responsible for ensuring that all tickets are resolved within the service levels that have been established for the organization.

The value of SPOC is that it brings order, discipline, and consistency to the support process. Support organizations that follow a SPOC model typically have lower costs and higher customer satisfaction that those that do not follow a SPOC process. Figure 7 on the following page illustrates a typical SPOC support process and the list below outlines key SPOC principles.

Key SPOC Principles

- The enterprise takes an end-to-end view of user support.
- The user/customer has a single point of contact for all IT-related incidents, questions, problems, and work requests.
- The L1 service desk is the SPOC and is responsible for:
  - » Ticket triage;
  - » Resolution at L1, if possible;
  - » Effective handoffs to other levels of support;
  - » Resolution coordination and facilitation; and
  - » Ticket closure.
- Desktop "drive-bys," "fly-bys," and "snags" are strongly discouraged.

Figure 7: SPOC Process

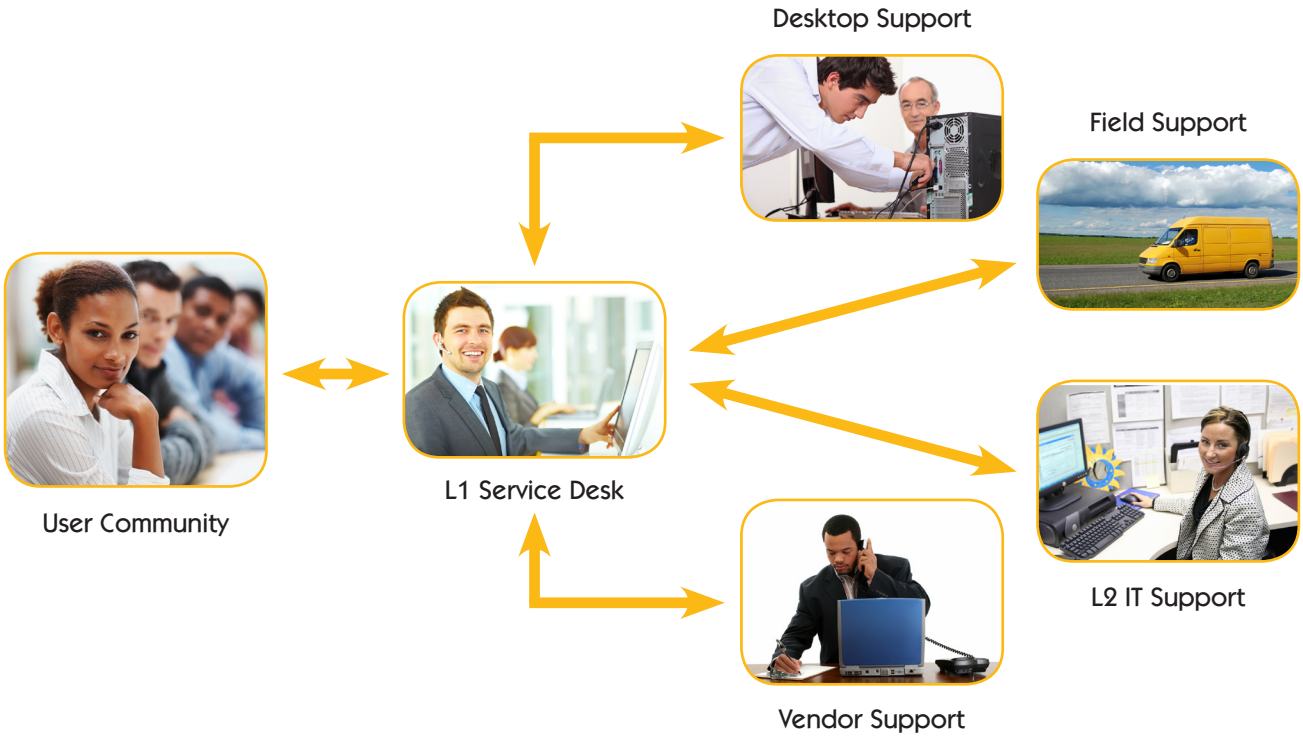

About the Author

Jeff Rumburg is a cofounder and managing partner at MetricNet, LLC. As a leading expert in benchmarking and re-engineering, Jeff authored a best-selling book on benchmarking, and has been retained as a benchmarking expert by such well-known companies as American Express, HP, and General Motors. He has more than twenty-two years of industry experience, much of it focused on desktop support.

Jeff received his MBA from the Harvard Business School, his MS *magna cum laude* in operations research from Stanford University, and his BS *magna cum laude* in mechanical engineering. Jeff can be reached by e-mail at **jeffr@metricnet.com**.

About MetricNet, LLC

MetricNet, LLC is the leading source of online benchmarks, scorecards, and performance metrics for corporate managers worldwide. MetricNet benchmarks focus on information technology, with an emphasis on desktop support, service desk, field services, and the call center.

MetricNet serves a global client base from its headquarters in McLean, VA. For more information, visit **www.metricnet.com**, e-mail us at **info@metricnet.com**, or call us at 703.992.8160.

About HDI

HDI is the world's largest IT service and technical support membership association and the industry's premier certification and training body. Guided by an international panel of industry experts and practitioners, HDI is the leading resource for help desk/support center emerging trends and best practices. HDI provides members with a vast repository of resources, networking opportunities, and the largest industry event, the HDI Annual Conference & Expo. Headquartered in Colorado Springs, CO, HDI offers training in multiple languages and countries. For more information, call 800.248.5667 or visit **www.ThinkHDI.com**.

