



The Exact Expertise to
Ignite and Sustain Performance™

LSA Global | White Paper



The Exact Expertise to
Ignite and Sustain Performance™

Skill-Based Routing

While skill-based routing has been around for a long time, it seems to be experiencing a sort of renaissance. It is probably fair to say that many of the early challenges have been overcome with new call routing and workforce management technologies, reporting tools, and just the learning that comes from experience.

Let's go back to the introduction of skills-based routing (SBR) and consider why it was embraced with such enthusiasm in the beginning. There are really four major reasons why call centers find SBR attractive:

- Gain economies of scale to improve service capabilities with lower staffing cost
- Minimize training costs
- Provide an orderly career development plan for agents
- Match up caller needs and agent capabilities

In many call centers there are a number of separate call types that are handled by separate groups of agents. For example, a group is dedicated to sales calls and another for service calls. However, the total number of agents needed to handle both call types together at the same service level may be substantially less than if handled separately – especially if the teams are small ones. Therefore, the “universal agent” or one who handles all call types is highly prized. Unfortunately, the cost to recruit and train these agents may be high or even prohibitive. Nevertheless, it is common for agents to possess the skills to handle more than one type of contact, even if every one can't handle them all.

It is SBR that allows the center to take advantage of utilizing every skill the agents have and gain as many economies of scale as possible while minimizing training cost. It also provides a platform for a career development plan that allows agents to start out with one skill and gain others and use them effectively as they learn. Trainees can be assured that they will not be given contacts to handle that are unfamiliar. Those who have just learned a new skill can be offered a few of those types of contacts interspersed with their more established skills to enhance the learning experience rather than being overwhelmed with call after call of the newest type.

For those call centers that have focused more on customer relationships and linking up the right caller to the right agent, SBR offers benefits as well. The benefits can be stronger customer relationships, reduced error rates, shorter handling time, and higher first contact resolution rates. Once again, with small volumes of work for specific customers or high-value groups, the economies of scale would be disrupted for dedicated agents serving just these segments. But with SBR, the agents can be focused on the high-value customers, while filling in “spare” time with other customer contacts.





The Exact Expertise to Ignite and Sustain Performance™

It is easy to see why call centers have readily embraced the concepts of SBR. However, not all the early implementations have been successful in delivering the expected results. Some of the initial challenges have included:

- Unmanageable complexity of design
- Limitations in workforce management (WFM) capabilities
- Limitation in reporting tools
- Agent resistance

Some center managers, vendors, and consultants got so carried away with the ability to sort contacts into specific groups, they sometimes ended up with nearly as many or more skills than agents. Every unique customer characteristic or agent skill level was defined in an effort to maximize the connection of right caller with right agent. This required the center to find a way to sort the contacts into these small “buckets” to ensure that they could be routed to the best agent or team.

Complex IVR scripts emerged that befuddled the callers, or many different 800-numbers were communicated to customers. However, customers just didn't comply with the plans. Some dialed “0” to reach an agent while others simply pretended to have rotary dial and waited for the default time out. Studies suggest that up to 25% of callers ended up in the default queue, and many of these had to be transferred to the “right” agents after answering. In addition, forecasting was essentially impossible when only a few calls per day were expected in some call types, and the half-hour in which they might arrive was any-man's guess.

The initial implementation of SBR was a challenge for most WFM vendors who had scheduling algorithms based primarily on variants of Erlang C. Queuing calls to more than one agent group, and agents logged into more than one skill simply defied the math of Erlang C. Complex algorithms evolved and simulation programs were added making WFM processes a closer match to the SBR routing plan.

The next challenge was reporting. After all, an agent logged into three skills might get a different mix of calls every hour, with varying average handle time (AHT) as well. A call initially thought to be a Sales call might be answered by the backup team in Service, making it difficult to determine what workload was being handled by whom. Vendors kept saying that it was impossible to report effectively on the fractions of agents represented by those logged into multiple skills.

For the agents, there was the conflict of the need to have coverage in all of the skills during all hours of operation and the rewards of seniority and performance. One problem is that when the call volume is low in each of several skills, a multi-skilled agent is far more valuable than a single-skilled agent. During the busy daytime hours, new hires with only one skill can be kept busy with just one call type, but at night and on weekends the multi-skilled agents are needed. This fact brings up the dilemma of rewarding someone for seniority or for learning an additional skill by putting them on the night shift. Some agents refused training, even when an increase in pay was involved. Trading a shift with a colleague also used to be easy and now the trade had to be between agents with the same mix of skills, limiting the possibility of finding someone to take an unwanted shift.

Another challenge was managing within the day when the actual volume of work and staff differed from the plan. Supervisors would access the ACD routing patterns and change the skills of agents or their priorities to try to get more staff into needed queues. However, they often forgot to change them back when the crisis passed. That left the ACD routing and the WFM assumptions out of synch so that any real-time reporting was totally skewed.





The Exact Expertise to Ignite and Sustain Performance™

It is fair to say that some call centers tried SBR and retreated to traditional routing, while others heard about the challenges and never gave it a try. Some implemented in a thoughtful step-by-step process and found the balance of complexity and manageability along the way. The good news is that we all learned much from both the successes and failures, and the chances of success today have been vastly improved by these pioneers

Here are a few of the lessons we learned:

- SBR can be a great boon if designed and managed carefully.
- The ratio of skills to agents is important.
- Forecasting requires minimal contact volume for accuracy.
- Schedule creation is complex and iterative, but doable.
- General shrinkage assumptions hold, but it matters “who” is missing more than ever – tracking schedule exceptions is critical.
- Utilizing the programming capabilities of the ACD can make intraday management much easier.

A Design Example

Skill-based routing certainly can contribute to the business needs of the contact center and enterprise. Let’s use an example from an actual contact center to demonstrate the possibilities and opportunities.

This center employs 42 agents and provides services in the health insurance field. Calls include those from both members (patients) and providers (physicians and hospitals). Some client organizations have negotiated service guarantees that the call center must meet or pay penalties, while others have no such guarantees. There are some government clients that generate very few calls but require that the count of calls and service provided be reported accurately. Most commercial clients are combined into a group where agents are cross-trained to handle all of them, resulting in much larger call volumes and better economies of scale. Analysis of the staffing required to meet the service goals suggests that the center could manage with 9 fewer staff if all were universal agents. While the center would like to have all agents fully cross-trained on all call types, that is simply not the case today and is not likely to be for months to come. New hire classes will always learn one skill to start, adding more as they progress. About half of the current staff are cross-trained in several of the 8 skills, and most possess at least 2 skills. Moving from a situation in which each agent takes only one kind of call at a time to a design that utilizes all of the possible skills of each agent is desired. This center does not process non-call contacts today but sees that potential in the near future. The table below provides an overview of the contacts and guarantees.

Skill	A Provider	B Com'l	C	D	E	F	G	H
ASA Guarantee	No	Yes	Yes	Yes	Yes	No	Yes	No
Call Volume per week	2300	2700	80	1800	80	50	1500	30
AHT (minutes)	8	7	7	7	7	7	4.2	4.2

The service guarantees are significant along with the need to manage queues with small call volumes to ensure they are not overwhelmed by those with large volumes. The center must be able to identify the source for each offered call, even if it is answered by someone who is a backup agent primarily assigned





The Exact Expertise to Ignite and Sustain Performance™

to another customer group. This is essential not only for the client reporting requirements, but also to maintain accurate statistics for forecasting and scheduling.

The first step in the process was to determine whether each of the 8 unique call types needed to remain in place, or whether some could be combined or even split. The contracts in place suggested that the 8 skills were needed for purposes of counting the calls and reporting. In addition, the 3 largest volume skills had a “backup queue” established (skills A, B and D). This allowed cross-trained agents to log into the backup skills to handle these calls, but maintained the integrity of the statistics so that a commercial call (B) was handled either in the primary commercial skill or in the commercial backup skill, but never routed to agents logged into other skills. It also offered the option of only using the agents in the backup skill if certain conditions were met as described below.

Determining the optimal number of skills is a crucial first step in the design. For every unique call type, a sorting method has to be put in place that can effectively separate these calls from the others. Will it be an IVR menu, a set of unique telephone numbers for customers to call, or a sophisticated CTI application tied to the calling line ID? In this case, a combination of IVR menus and separate telephone numbers was already in place and working adequately. It is important to be aware of the potential for some number of callers to make the wrong choices or refuse to utilize the sorting system. These will end up in a “human triage” situation and if it is a significant percentage of the calls, it can be a major design consideration. These calls must be added to the volume normally handled by the default agent team, and may result in a large percentage of transfers.

Next, the relative importance of certain callers over other callers was determined. This was based on the service level guarantees in this case, but could be a function of customer value or other business drivers. In this case, all calls were considered equally important, at least as they entered the queues. As the service levels declined in the guaranteed groups, the priority of these calls would be raised above the other calls to ensure they were answered quickly.

In the next step, the skills and relative capabilities of the agents was considered. A matrix was constructed with all of the possible skills and relative skill in each for every agent. While the ACD offered over 40 possible levels of skill to differentiate agents from one another, the initial design should keep it fairly simple and in this case only three levels were used. These skill levels were assigned to the call types for each agent so that if several agents were available, the next call would be routed to the agent with the best matching capabilities.

The table below defines the agent skills and the number of agents having each skill. In each team the center defined some agents as primary and others as second or third in deciding which should get a call of that type if several agents were available. For example, there are 18 agents trained in the B team and 5 of them are considered primary to that team while 13 are secondary. Notice that the 5 on the B team can also handle calls for the A and C calls, but do not handle the D calls, with the other 13 can do all of these. It is also apparent that some call types have no primary agents and are totally reliant upon agents who are primary in other skills (i.e., C, E, F, and H). These skills exist more for counting and reporting than for staffing purposes. Each will require a separate forecast of workload in the WFM process but will be combined with others who share the agents when the scheduling is done.





The Exact Expertise to Ignite and Sustain Performance™

Skill Type	A	A BACKUP	B	B BACKUP	C	D	D BACKUP	E	F	G	H
1ST choice	A 8	B 5	B 5	A 8	D 1	D 6	B 5	D 2	D 6	G 2	G S
2nd choice		B 13	B 13		B 5	D 2	G 6	D 6	D 2	G 6	G 6
3rd choice		D 2	D 8		B 13	B 13		G 6		D 2	

The next step was to consider the concerns of small call volume clients (with service guarantees), and the balancing of agent utilization to ensure that the most important calls were handled before the least important. This required the use of call priorities with ASA drivers. The “if-then” scripting statements in the ACD were used to check the ASA (over the last 10 minutes) in the small queues. If the ASA started to slip below a certain level, then the next call was bumped up to a higher priority than all other calls in queue. This ensures that this call is answered quickly, but since there are likely to be only one or two calls per period, the impact on the other call types’ service levels should be minimal.

In the larger queues, ASA checks also could be used to determine if a call should be prioritized based on the service guarantees. Where appropriate, priorities could be raised as the ASA slipped below an acceptable level, but this could significantly impact the other types of calls these agents were able to handle. In this case, the “if-then” scripting statements checked the ASA to determine when to tap into the cross-trained agents who log into the “backup” skill. So the primary agents are utilized as frequently as possible, but backups can help out when needed.

The result of this design should include a relatively even workload among the agents throughout the center, filling in idle time with backup call work. The service should also be somewhat more consistent among the groups and hours of the day. However, there is at least one group that is considered “last priority” to the center (Skill A) and could see its staff utilized primarily for Skill B calls. Skill A is considered the “sacrificial lamb” in this center and it could see a poor service level if Skill B calls overwhelm the available resources, but that is acceptable in this company. No agent team is “an island” with all of the backups and cross-skilling, so the schedule for this entire office must be done as in a single process to effectively consider the utilization of all personnel on all skills.

In calculating the total resources needed in this center, it is difficult to determine the exact head count required without a sophisticated simulation, but some gross approximations will help demonstrate the benefit of economies of scale achieved. If each of the 8 call types is staffed with a dedicated team, the total bodies in chairs requirement in the busy hour is 45. In the average hour, the total requirement drops to 36 in separate teams. With a current staff of 42 (all working 8-5 M-F), and an anticipated shrinkage of 20%, the available staff is closer to 34 during most days.

If all agents are able to handle all calls, then the total bodies in chairs required in the busy hour drops from 45 to 37 and in the average hour from 36 to 28. This center has not yet achieved the full cross-training to utilize every agent for every call, but with this SBR design, they are able to achieve a result that is somewhere between the separate skill requirement and the universal agent requirement—where on the scale will vary by day depending on which agents are absent and by half-hour as the call workload shifts from one call type to another.

Those skills with less than 100 contacts per week are a forecasting nightmare. It is impossible to say in which half-hour today’s 10-25 calls might show up. So the center must simply ensure that there is at least





The Exact Expertise to Ignite and Sustain Performance™

one agent logged in for each of these at all times, even if the standard scheduling process might suggest otherwise.

Once the ACD scripting was implemented, the call center saw rather dramatic improvement in both customer service levels and even workload distribution among the staff. By gaining some additional staffing in the busy skills, service was improved. Filling in idle time for agents in low volume skills helps even out the workload for all staff. Even the “sacrificial lamb” skill (Skill A) saw an improvement in their speed of answer from 9 minutes to less than 1 minute.

In fairness, this example call center didn't have much of a scheduling problem due to the 8am-5pm, Monday through Friday hours of operation. But for those with longer hours, part-time staff, and more complex work rules, the WFM process will need to be able to replicate the routing environment for this to be effective.

Sorting the Contacts

While it is clear that sorting the call and contacts into the various skill groupings is a key element of the SBR design, the good news is that the technologies for accomplishing it have improved. Speech recognition has made it possible to use automated menus in more complex environments and increase the percentage of callers who end up in the right agent team (if not processed to completion by the automated system). Customers are getting used to the process of menus and sorting and seem to see the benefits of ending up with someone who can actually resolve their problems. It is important to design based on the way a customer thinks when sorting the contacts, not on the way the center would like to utilize its agents. If the customers cannot figure out the right choice, or simply choose not to participate, the design will fail.

As the automated contact handling systems have developed, they are handling more and more contacts to completion without the aid of a human in the call center. While this is a cost-effective solution, it has raised the average level of complexity of agent-handled contacts. With most of the simple and quick transaction calls handled in the systems, only the most challenging problems are referred to the agents. This increases the AHT on the remaining calls and suggests that high-caliber reference tools and knowledge base products are critical to achieving first call resolution and high quality. We must also consider the potential to mix calls with written contacts such as email and web chats. While the product knowledge may be the same for both, the basic skills of handling calls and writing sentences are not. Developing agent skills in written work is a critical element of inter-mingling these two types of work, and many call centers have kept the written and verbal work largely separate as a result. Once again SBR can provide the opportunity to have some agents dedicated to one or the other, while a few do both. The reality is that making significant gains on the staffing efficiencies does not require that all agents be cross-trained and universal. Just a few can make a huge difference by crossing over when the workload in one or the other changes.

Routing the Contacts

Once the contact types have been sorted out, routing them to the right resource is the next step. Great strides have been made in routing technologies so that even multi-site teams can be effectively utilized as if they were in one building. Calls can be routed through delivery to one site and continuous testing of other site situations to determine if there is a shorter queue elsewhere. This might be effective when the IVR is on site and even when only at some sites in the network. Or calls can be routed to the “best” site initially by a routing processor that communicates with the public network and each site's ACD to determine where the call should be handled. VoIP systems offer the option of dealing with all agents at all sites as if they were in the same room. Even in a single site, calls can be queued at the network rather than one site if desired.





The Exact Expertise to Ignite and Sustain Performance™

The challenge here has often been maintaining visibility of the process and where the work is being routed and handled—but even that challenge is being addressed by new real-time reporting tools. Views can be seen of the service levels, volumes, AHTs, etc. on a per-site, per skill basis so that appropriate adjustments can be made before the situation deteriorates.

Creating the Schedules

When it comes time to match up the agents to the needed bodies in chairs for each skill, the challenge is significant. Creating a schedule that utilizes multi-skilled personnel goes far beyond the capabilities of the typical WFM Erlang C tools. Each person will be utilized part of the time in each skill defined. If the schedule is being done for a week when three agents with single skills are on vacation, the plan will be somewhat different from the week when the three agents on vacation normally handle a combination of 5 skills. It is still 3 people out, but which agents make a great deal of difference in how the schedules will need to be constructed and assigned. Simulation is needed to try a variety of possibilities to determine the “best” schedule that can be made up from the agents that are assumed to be available.

Then reality sets in. Several agents call in sick or fail to adhere to their schedules. Again, it really matters which ones they are just as it mattered when the schedule was designed. Agents will want to trade shifts, but care must be exercised to ensure that either the two agents have the same skills or the impact of the change is not unexpected or significant to the results.

This is where the power of the ACD adjustments can really help. If the ACD is programmed to check conditions on every call, it will react as directed only when that condition exists. For example, if there are more calls in queue than agents logged in, it might make sense to tap a backup team. When that happens, the next call will go to the backups, but when the next call arrives, the ACD will check to see if the condition still exists before sending it to the backup as well. In this way, the calls will be delivered to the teams that can best manage the work without any human intervention both during and after the crisis is over.

Developments in WFM Processes

As mentioned earlier, WFM tools initially struggled to keep up with the capabilities of the ACD and routing processors in SBR. This utilization of agents among a variety of queues has defied utilization of the standard queuing algorithms. How can a call center manager determine how many people to have on the payroll, when to train more in a particular skill, and how to grant time-off requests when the number of staff needed in each call type is constantly in flux? Development of the staffing plan and schedules is dependent on being able to establish a set of assumptions that closely replicate the SBR routing rules and going through an iterative process of trying out a variety of options until the optimal solution is found—not something that can be done in a spreadsheet or Erlang C model.

The management of schedule exceptions (making notations of who is not available as planned) is far more critical in SBR than a traditional dedicated agent environment. The key is who is missing and which combination of skills does that person have? If a single skilled agent is missing, it may create a big gap in one team, but it will also affect the teams who have cross-trained agents who will spend more time today covering for the loss of that person. When a multi-skilled agent is missing, the impacts may be felt throughout the organization and several skills are impacted and coverage must be adjusted. In small call-volume skills, one agent may be all that is scheduled during some periods, and these customers may go into infinite queues if adjustments are not effective.

To make matters even more challenging, some of the work is randomly arriving such as the inbound phone calls, but some is sequential including the emails and correspondence. Scheduling agents who can move back and forth between different media requires a WFM system that can handle both random workload and the sequential kind that builds up a backlog. In addition, the WFM tool will need to be able





The Exact Expertise to Ignite and Sustain Performance™

to accept input from a variety of sources including the standard ACD link, link to a routing processor, and links to other systems where work is managed, such as an email or web chat management system.

Once the system is fully set up to replicate the variety of contacts, arrival rates, and staffing options, it needs to stay synchronized with the routing tools. This means that changes to the skills and priorities that operations supervisors tend to make during each day to react to changing needs must be reflected in the assumptions in the WFM system. This may mean changing the process to only allow WFM to make such changes in the systems, or a better communications process to ensure everyone is making the same assumptions as changes to the schedule are discussed. Better yet would be reliance on the scripting in the ACDs themselves to manage the moment-to-moment adjustments since they react faster than humans and always reset to the default when the crisis is past.

Checklist for Effective SBR Design

Developing and managing a SBR design is complex with a number of processes each affecting others. A checklist of issues to be considered and managed is provided below to assist in both the initial design and the consistent management of these types of centers.

1. Identify the business needs to be addressed

- a. Are some calls more important than others?
- b. Is training cost for universal agents too high or unrealistic?
- c. Is there a need for a progressive career path for agents?
- d. Is the workload unevenly distributed among the staff?
- e. Is the service level inconsistent between skills and times of day?
- f. Are there some low volume contact types that need protection from being overwhelmed by the high volume contacts?

2. Analyze the contact sorting process for effectiveness

- a. When callers opt out, do they route to the right agents?
- b. Does the screen pop the information from the automated input?
- c. Does the agent effectively use the screen pop data to minimize AHT and reduce asking customers to repeat information already given?
- d. Is the email sorting process routing effectively to the right agents?

3. Agent utilization plans

- a. Does the design balance the skill coverage needs with agent needs?
- b. Have procedures been put in place for handling the shift bids, trades, and time off requests effectively?
- c. Has the center built a training and career development plan around skill coverage needs?

4. Maximizing the effective utilization of WFM technology

- a. Have small volume skills been combined where possible for forecasting accuracy?
- b. Is there an effective methodology for developing long-term hiring and budgeting plans?
- c. Does the WFM process include simulations in the scheduling cycle?
- d. Are effective tools and processes in place to manage within the day?
- e. Does the center work effectively with the WFM vendor to ensure future needs are met?
- f. Is there an effective plan for training and development of the WFM team?





The Exact Expertise to
Ignite and Sustain Performance™

Summary

It is fair to say the skill-based routing has changed the world of call centers forever. The opportunities to improve service and efficiency at the same time are few and that is the promise of SBR. The early implementations often struggled to find the right staffing design to match up to the routing design, but the WFM tools and other supporting technologies have met the challenges and are ready to make it possible to achieve solid results. If SBR would help to meet some of the business challenges in your call center, it is time to give it another try and realize the benefits others are achieving.

About LSA Global

Since 1995, LSA has helped organizations create and maintain competitive advantages through people. Over 85% of our business comes from repeat clients and our customer satisfaction rating is 97.5%. Our clients tell us that we are different. Our clients tell us that we save them both time and money. Our clients tell us that they appreciate access to experts across many areas

Our Leadership and Management clients report that we decrease their costs by up to 50%, while helping increase speed-to-productivity by 60% and decrease unwanted attrition by up to 40%.

Our Sales clients confirm that they have grown sales by 40%, increased units sold by 42%, increased average pricing by 12%, and closed over 50% of their pipeline.

Our Project Management clients inform us that, by regularly completing projects that meet or exceed expectations, they have become an anomaly in the world where most projects disappoint or fail.

Our clients succeed in the marketplace through increased revenue, decreased costs, and higher productivity. They maintain that our rigorous assessment, implementation, and measurement capabilities bring them tangible results. As an organization, we are fiercely devoted to supporting their success.



Call Us Toll-Free 800.889.6452

Copyright ©2007 LSA Global. All Rights Reserved. All copyrights, trademarks, service marks and product names are copyrights, trademarks, service marks, or registered trademarks of LSA Global or its Partners.