



Split Decision:

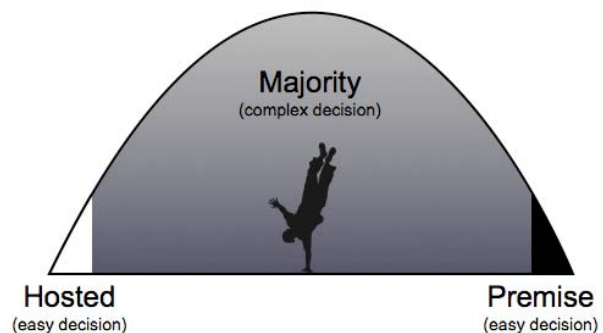
How Vendor Bias Affects the Complex Hosted vs. Premise IVR Business Case

**Nancy Jamison
Jamison Consulting
July 2007**

Premise or Hosting: Tough Call

If you have decided to add IVR – particularly speech-enabled IVR - to your customer service arsenal, you now face a decision that could have a bigger impact on your business than the IVR application itself. Should you deploy it on your premise systems or with a hosted IVR service? Prevailing arguments from biased vendors make the decision seem easy: go premise if you want ownership and control; go hosted if you want flexibility and minimal capital expense. Despite what you may hear, it's just not that simple - and making the wrong decision could cost you dearly.

It turns out that the tidy arguments from vendors are black-and-white edge cases that clearly favor the hosted or premise solution they want to sell you. If you happen to be one of the few customers who live on the edge where these easy business cases apply, good for you. But if you're like everyone else, you fall into a huge gray area where the decision to go premise or hosted is much more complicated than it first seems. In this sea of gray, many buyers fail to consider key aspects of this complex business case.



This whitepaper examines some of the hidden forces that make that gray area so gray. It will help you navigate through the gray area by helping you frame how those forces might affect your company. It will also help you clear some of the fog by suggesting ways to bring those forces under control.

Premise and Hosting: Both Mainstream

At one time the entire IVR market was black and white. Hosting was purely an edge case that made sense for a few, and everyone else bought premise. While I was at Gartner in 1995 the total IVR market size was \$1.006 billion, and only 8% of that was hosting. I forecast the IVR market to grow to \$2.7 billion by 2000, with hosted services growing to 9% of that total. At the time, part of my reasoning for curbing hosting growth was that only a handful of vendors were providing hosted IVR services, and nothing new was happening that would make hosting appeal to a broader set of customers.

In 2007 some research reports find premise and hosted IVR running neck and neck for the leading share of IVR market revenue. Times certainly have changed. Lower transport costs

and the adoption of VoIP with SIP leveled the playing field and improved the economics of hosting in general. Widespread adoption of VoiceXML has given enterprises unprecedented control over how hosted applications are created, deployed, and modified compared to the proprietary hosted IVR architectures of the past. As expected, this market opportunity has attracted many new IVR hosting vendors, sparking a hosted vs. premise market battle that continues to be a hot topic in the press and at every industry gathering. The message is clear: hosting has become mainstream.

With both sides strong in the market and neither dominant, it's harder than ever for a decision maker to run the gauntlet through these arguments and decide whether to deploy on premise or host. Let's review the basic arguments for each before we consider the hidden forces that further complicate matters.

The Basic Arguments

Premise-centric vendors base their arguments on lower total cost of ownership (TCO) and more operational control. Though premise deployments can require a significant up front investment, per-minute fees charged by hosting vendors add up quickly. Lower operational costs for premise deployments eventually offset the higher capital cost, and premise systems become less expensive over the lifetime of the deployment.

Giving up total operational control over their systems and applications is not easy for most people. *Premise-centric vendors* know this. They argue that premise deployments allow you to maintain control of system uptime, quality of service, and maintenance windows. Your own priorities and resources will drive the implementation of new technology and allow you to push the envelope with unique features. You'll also avoid depending on another party to keep your data and equipment secure.


Hosting-centric vendors argue that having no up-front capital cost and the ability to deploy rapidly make hosting more desirable than premise. Vendor-supplied infrastructure, maintenance, technology upgrades, and domain expertise are all rolled into operating expenses. Because the hosting vendor spreads these costs over multiple customers, individual hosting customers enjoy economies of scale that give them a higher level of service at lower cost than building and maintaining a similar infrastructure and operations capability internally.

It's expensive to build and maintain reserve capacity on premise, and once capacity is deployed it's not easy to reduce costs if that capacity isn't being used. *Hosting-centric vendors* know this. They argue that on-demand scalability allows you to scale hosted deployment costs up and down as needed so you can more easily deal with seasonal demands or unexpected changes in call volume.

What's Wrong with These Arguments?

These arguments describe clear-cut cases at each end of the premise-hosting deployment spectrum. The problem is this: practically nobody fits the edge cases. Almost everyone falls into the huge gray area in between.

Business cases based on these arguments hang in the balance and can tilt quickly one way or the other as assumptions change. In fact, failure to properly examine two factors - risk and change - can cripple the deployed ROI of an IVR project that once looked good on paper. Biased vendors don't communicate the importance of risk and change because it introduces enough complexity and abstraction to crush an elevator pitch. Oversimplifying these arguments makes them easy to understand, but they aren't much help to most customers struggling with the decision to deploy on premise or host. Instead of dwelling at the edges, let's concentrate on how to navigate through the huge gray area. To do this, it's important to understand how risk and change make that gray area so gray.



Risk

In order to build a business case to help you decide whether to deploy a premise or hosted solution, you're going to have to make some assumptions. Since predicting the future is *really hard*, chances are some of your assumptions will turn out to be *wrong*. No shocker there. But here's what turns a black-and-white business case into a gray one: surprisingly small differences between assumptions and reality can tank your ROI so badly you'll wish you had made a different deployment decision. Worse yet, a typical cursory sensitivity analysis probably won't catch this.

Let's consider just one key risk, but potentially the most difficult to grasp fully: **traffic risk**. When call volumes are steady (not much change between the day-to-day peaks and valleys) this isn't a big problem. You'll know roughly how many minutes you will need to purchase from a hosting vendor each month, and based on that it's fairly easy to calculate the cost to deploy enough capacity to handle an equivalent amount of traffic on premise. You might even do a little sensitivity analysis and look at how a 20% increase or decrease in monthly call volume would impact your costs, then choose the deployment strategy that looks more favorable overall.



What happens if that same call volume turns out to be *not quite as steady as you thought*? Let's say instead of being spread fairly evenly throughout the day, traffic turns out to cluster around a certain time of day - a phone caller rush hour of sorts. Just like highways always

seem to need more lanes for cars at rush hour, your IVR system will need enough capacity to handle call volume peaks if you want to keep processing calls smoothly during the rush hour. So the sensitive question is *not* how many minutes a month of traffic you expect, it's how many minutes of traffic will jam your system during rush hour.

How does this traffic risk affect your business case? If you assumed a fairly even call distribution throughout the month and later discovered that you had a large daily rush hour peak, your business case could be in trouble. Depending on the height of that peak, your infrastructure cost to deploy enough capacity on premise *could easily double* compared to your original assumption - even though the total amount of monthly traffic *did not change*.

Similar traffic risk problems could arise from an unpredictable burst of traffic. Three hundred fifty days of the year your assumptions are spot on, and then whammo! You get hit with a traffic spike that redefines your worst-case scenario. You're a utility and a storm-of-the-century interrupts service to the whole city, and your inbound call volume skyrockets in less than thirty minutes. Or the marketing department of your financial institution does a retail product promotion that suddenly catches fire. Interested customers jam your phone lines and you shudder to think how many sales you are losing because customers can't get through.

You might think that fear of traffic risk would always steer you from premise to hosted so you'll have on-demand capacity to handle the spikes, but it can easily work in both directions. Your assumptions about the amount of peak capacity you require will be a big cost driver in your business case. Assuming too high a peak could drive you to deploy a hosted solution because deploying enough capacity on premise would appear prohibitively expensive. But if the actual traffic later turns out to be more predictable than you had feared, you might have been able to get away with a more modest investment in premise capacity - which could quickly flip the business case and leave you wishing you had deployed on premise to achieve a better ROI.



"You *should* have bought a *smaller* glass!"

Change

So you've made your assumptions, analyzed all conceivable risk and finally feel like you've got it all figured out. You choose a strategy and deploy. Sure enough, a change happens that rocks the foundations of your carefully constructed business case. Your ROI is suddenly underwater as your initial assumptions evaporate overhead. Change happens.

Change is indeed inevitable, and it can come from anywhere: a corporate merger that shuffles priorities or available infrastructure, a geopolitical event, new fad, or disruptive technology that triggers a radical shift your traffic profile, or economic conditions that wipe out your capital spending budget just when you need to buy more capacity. How can you possibly deal with something so unpredictable? Most people ignore the possibility of change because they can't predict it. That's not the best way to deal with something you *know* is inevitable.

Vulnerability to change helps expand the huge gray area precisely because most people aren't comfortable dealing with it. If you want to be a superstar, embrace change. When you build and analyze your business case, *consider how easily you can respond to change* - any change - especially if it requires a radical shift in your deployment strategy. The more easily your deployment strategy allows you to respond to change, the better off you are. In fact, it's much better and safer to invest in making your deployment strategy more tolerant of change than it is to invest in trying to predict the future more accurately so you can choose a deployment strategy that's optimized for your prediction. Change is inevitable, and predicting the future is really hard. Optimize for change, not predictions.

The more easily your deployment strategy allows you to respond to change, the better off you are.

Protection from Risk and Change

With hosting now in the spotlight, the burning question of whether to deploy a premise or hosted IVR solution remains one of the hottest topics in the industry. Yet there's an even better question that's so obvious many people forget to ask: Why not both?

Premise or hosted... Why not both?

Premise and hosted deployments each have their advantages, but a combination of the two can be even more powerful than choosing one or the other. If done right, this strategy will let you enjoy the best of both worlds and make your business results considerably less vulnerable to risk and change.

Scenario 1: Crystal Ball Strategy

Predicting the future is really hard. Wouldn't it be great if you had a chance to test your assumptions before you had to commit to a rollout deployment strategy? Using a hosted system for trial or initial deployment is a great way to collect traffic data as well as tune the application to ensure it meets your objectives. Armed with this data, you'll be in a much better position to analyze your business case and determine your best long-term deployment strategy.

The Crystal Ball strategy also works well if you need to prove the benefits of a new application before launching on a larger scale. If the application flops during a hosted deployment, you won't have to explain a large capital investment in premise equipment

that's no longer necessary. If it succeeds, you'll have traffic data from the pilot and quantifiable data on application benefits you can use to secure funding for a larger rollout—perhaps this time on premise.

How does this protect you from risk and change? It not only helps reduce forecast risk in the short term, but also allows you to establish a relationship and successful deployment with a hosting vendor that could be valuable should your circumstances dictate a need for hosted capacity in the future.

Scenario 2: Rapid Deployment Strategy

Sometimes you need to get an application deployed in a hurry, but circumstances make it difficult to deploy quickly on premise - maybe you need more time to build out infrastructure or perhaps you don't have capital budget until next year. You can launch your application with a hosted IVR service, then migrate to premise as soon as you are able.

The Rapid Deployment strategy protects you from risk and change by giving you considerable flexibility on timing the migration. Should any changes suddenly rattle your business case and delay your premise deployment, you're already covered by a successful hosted deployment.

Scenario 3: Cloudburst Strategy

If you deploy both on premise and with a hosted service at the same time, you can split capacity between the two in ways that let you enjoy the benefits of each. While premise can be lower cost if you have existing infrastructure to leverage, hosting is particularly good at handling bursts in traffic.

For example, you could deploy enough capacity on premise to handle most or all of your baseline traffic, and combine it with supplemental hosting capacity "in the cloud" to handle all of your bursts in call volume. Though you would still need to pay a hosting vendor a fee to keep enough peak capacity in reserve, this would be offset by your need to purchase less premise capacity. Instead of having to put enough capacity on premise to handle your busiest-day-of-the-year peak call volume, you could buy just enough capacity to handle your average daily call volume and let the rest burst into the hosting cloud.

The Cloudburst strategy protects you from risk and change by giving you a tremendous ability to roll with the punches. With hosted and premise deployments operating in concert you can periodically (and quickly) respond to almost any disruptive change by adjusting the mix of traffic handled by hosted and premise systems. Hosting protects you from the traffic risk of large unexpected bursts in call volume, and you can conservatively ratchet up premise capacity only as fast as you're absolutely sure you have enough traffic to keep it fully utilized. Furthermore, this gives you additional flexibility to manage capital and expense budgets by adjusting the ratio of the hosting/premise split over time.

Scenario 4: Hosted Redundancy Strategy

Maybe you favor an on-premise deployment, but your application demands high system availability. The expense of deploying the necessary redundant, geographically dispersed

infrastructure and hiring the expertise to keep it running would be cost-prohibitive. Your IVR project looks doomed unless you can come up with a strategy that radically improves your business case. Hosted redundancy might be a perfect solution. You can deploy your primary system on premise, but arrange for failover traffic to be handled by a hosted service. This provides redundancy and geographic dispersion at a fraction of the cost of building it yourself. In fact, the economics of this approach can be so enticing that hosted redundancy makes a failover solution feasible for not only the most mission-critical applications but also those applications where high availability is merely desirable.

How does this protect you from risk and change? There aren't many changes more disruptive than complete system failure, and hosted redundancy helps you maintain business continuity in the event of what would otherwise have been a disaster. In addition, a hosted redundancy deployment poised to take failover traffic can easily be converted to a Cloudburst arrangement that handles overflow traffic as well.

Making a Premise + Hosted Strategy Work

Complementary premise and hosted deployments can be extremely attractive if executed properly. As a general rule, these strategies work best when you can make them look more like a single deployment with premise and hosted components rather than separate deployments that each operate in a vacuum. Be sure to look for the following:

Platform compatibility. If your premise and hosted deployments use platforms from different vendors, or even different versions of a platform from the same vendor, annoying problems could arise. You want your application to run seamlessly and with identical features in both places, so you don't want different platforms to cause you to have to maintain two separate versions of your applications and keep them in sync as they evolve. Luckily, industry standards such as VoiceXML have reduced the magnitude of this problem in recent years, but don't be lulled into a false sense of security. Implementations of standards still vary, and even certified VoiceXML compliant platforms vary in the number of "optional" certification tests they support- so make sure you insist on a robust and complete implementation of the standard that passes all required and optional certification tests if you want maximum flexibility.

Contract compatibility. Part of the value of a "protection from risk and change" deployment strategy is the ability to make adjustments over time that reflect your changing circumstances. If your premise system and hosted service contracts are too restrictive and don't allow for this, you're back to square one. A single contract that covers both is ideal, but separate contracts will do if you retain a degree of flexibility. Should you need to increase or decrease your commitment in order to shift capacity from hosting to premise or vice versa, you don't want to have the benefit wiped out by residual commitments for capacity you don't need.

Excellent Customer Service and Support. This should be obvious, but if you're going to take advantage of a deployment with premise and hosted systems running in parallel, you're going to want to make sure you select vendors with excellent customer service and well-

trained support. To take full advantage of this strategy you will want to ensure that failover and overflow plans are properly established and provisioned so that they can be implemented quickly and easily. The better support you receive from your vendors, the better command you will have over the combined deployment- especially when you need to make adjustments.

Conclusion

The business case to compare a premise vs. hosting deployment strategy is very complex. Most vendors oversimplify arguments in favor of hosting or premise by describing clear-cut edge cases that fit only a small fraction of customers. In reality, most customers trying to make a decision find themselves in the huge gray area between the obvious hosted or premise edge cases.

To deal with the uncertainty in the gray area, most decision makers struggle to make increasingly accurate predictions of future costs, traffic, and benefits in order to drive their business cases to a clear choice of deployment strategy. This effort is futile, because the effects of risk and change almost always outweigh one's ability to accurately predict the future.

Instead of analyzing whether to deploy hosted OR premise IVR, consider how to leverage both hosted AND premise IVR in complementary ways. Deployment strategies that incorporate complementary hosted and premise systems can provide significant protection from risk and change if implemented properly. The key to proper implementation is twofold: 1) make the overall solution behave as much as possible like a single deployment with premise and hosted components, and 2) make sure you have the ability to control and adjust the mix of traffic between the premise and hosted components at any time. Doing so will allow you to maximize your "deployed" ROI by responding quickly to anticipated risks or unanticipated changes that might otherwise undermine the predicted ROI for a less resilient premise-only or hosted-only deployment strategy.