

Under the Buzz

Commentary on Business Strategy Issues for Executives in Enterprise Systems & Software Companies

July 31, 2006 - Vol. 7, Number 3

Under the Buzz is an email "newsletter" authored by Philip Lay, managing director at TCG Advisors, a Silicon Valley-based firm that helps to *catalyze* the innovation and transformation efforts of executive teams in systems and software companies. This journal is published periodically and delivered free to subscribers via email. It is also posted on TCG Advisors' website (<http://www.tcg-advisors.com/Library/utb/utb.htm>), back issues are also available.

Topics in this issue:

1. **The Problem with "Bubble-Up" Innovation**
2. **Case Study: Why it is Vital to Understand Customer Resistance to New Innovations**
3. **Dealing with The "Black Box" Syndrome and "Meeting Hell"**

1. The Problem with "Bubble-Up" Innovation

It is a more or less a given in today's globalized marketplace that every company needs to innovate consistently in order to differentiate itself against its 'competitors' – whether these be other product or service providers, or prevailing ways of doing things (the proverbial 'status quo'). Tech companies, needless to say, are no exception. In fact, consumers and corporate users alike almost certainly feel disappointed when tech companies fail to positively exude innovative energy, particularly when they have their investor hats on.

But are tech organizations really that good at innovating? If you just consider the amounts of funding and resources they profess to direct towards development and engineering – and R&D budgets typically range from 8% to as much as 24% of revenues – you'd probably think that (a) they take it very seriously, and (b) they know what they're doing.

Well, not so fast. Software and systems companies get it wrong a lot of the time. For one, they tend to encourage innovation to bubble up from anywhere in the company, possibly due to their entrepreneurial origins. But even staid old Fortune 500 companies such as 3M and BASF have gained attention in recent years for their ability to generate revenue growth largely from recently launched product innovations. In general, it's been fashionable for many years now for management to democratize innovation, rewarding employees with bonuses for creative ideas that lead to new products and services. On a purely philosophical basis, it's difficult to object to the notion of empowering employees and even partners, and others, to contribute to a company's success in the marketplace. It is highly motivating, can be extremely rewarding, and everyone – including of course customers - can share in the successful outcomes.

In tech, however, the problems start because company executives allow "bubble-up" innovation projects to go *unmanaged*. In fact, it's not infrequent that the executive team as a whole knows little or nothing about a given new alpha product until resource requirements beyond the authority of the "innovating" individuals or teams become necessary. At which point, the opportunity cost – and even the direct cost – of having allowed those individuals to scope and get the project moving may be quite considerable.

Perhaps not surprisingly, this bubble-up problem is especially prevalent in companies that lack clear focus and differentiation in the marketplace. Often, these organizations are rife with unconnected one-off projects. Usually, this trend is driven by an increasing anxiety among executives and employees alike regarding the company's relevance as a player in their target markets. I see this phenomenon as an acute and costly form of organizational *fidgiting*.

Another symptom in the industry that my colleague Geoffrey Moore has commented in recent articles and speeches on is the habitual hyperbole surrounding R&D budgets that is spoken by company executives at industry trade shows and analyst briefings. Their gist tends to be a variation on something like this: “We are spending \$x millions/billions every year on innovation”. The clear implication being, “*That’s* how innovative we are!” The trouble is, for most hardened listeners it comes across as exactly what it is, i.e., rhetorical one-upmanship (as in “My R&D budget is bigger than your R&D budget”). Ironically, it is the CEOs of companies with the poorest track records of *effective* innovation – in other words, innovative products that actually garnered success in the marketplace – that make most noise in this direction. After all, any metrics less than revenue growth from said innovations, or – even better – stock price appreciation resulting from the new market power achieved from the innovations (Apple’s iPod/iTunes offering being everyone’s current poster child) are hardly evidence of anything except possibly “We try harder,” or as skeptics might say, “We spend faster,” or as cynics might say, “We waste faster.”

Over time, we have come to understand that the secret to understanding the anomaly of companies – including household names such as IBM, HP, Intel, Oracle, and Microsoft, to name but a few – that repeatedly claim to be spending big bucks on technology-related development but have little or nothing to show for it relates at least in part to the simplistic way they choose to define and measure “innovation.”

The Definition Problem: Innovation is defined narrowly, as developing something “new.” But new to whom? And new what? What if, as is so often the case, “new” means nothing more than a feature or an offering that the company itself has not been capable of offering before? If one or more competitors have already developed – and even launched – an equivalent feature, product, or service – how innovative is the first company’s innovation now? Furthermore, what counts as innovation?

To avoid continued oversimplification in this regard, our firm has gone as far as to conduct extended field research into this topic. The result of our efforts is a taxonomy that describes 14 different innovation vectors, or forms of innovation, that tech and other companies must at least consider before they decide on their innovation strategy(*). Our hypothesis is that, before authorizing any ‘bubble-up’ innovation efforts, a company must do sufficient thinking about their specific innovation strategy in order to reduce the 14 possible choices to at most one ‘major’ and one ‘minor’ vector. For example, a company could decide to major in enhancement or design innovation and minor in value engineering or integration innovation, each of which implies quite different investments. To complete this same example, the company would implicitly not invest in line extension, marketing, experiential, process, or other forms of innovation, based on the assessment that these would not be rewarded with price premiums or easier access to their target markets. Management can then publish this innovation strategy to their employees and partners in order to frame the type of things that they will treat as priorities for innovation investment, against those vectors that they cannot see any future reward for investing in.

At the end of the day, innovation is hard, risky work, not to be undertaken lightly, nor – and this is a key point for many techies – for its own sake. Let’s not forget that the purpose of innovation, at least for for-profit companies, is to gain competitive separation against all known *alternatives*, be they other categories, other markets, other products, or other companies.

The Measurement Problem: We’ve already discussed the input side of the equation related to and other investments in attempted innovation. What about the outputs or results? In our experience, perhaps the most frequent misleading ‘results’ metric that drives R&D behavior – and bonuses – is the number of patents registered or applied for. When companies set out to prove their innovation supremacy by citing the number of patents registered or awaiting approval every year, they usually fail to impress anyone, save perhaps their own PR staffs. Software patents in particular are very hard to protect, and competitive advantage can be fleeting, at least in comparison to the time it takes bureaucrats and lawyers to complete the patent process.

The second most common misleading metric is probably the number of features or products taken from concept to beta by R&D teams (or something to this effect). This often leads to finger-pointing between R&D and marketing or sales, when they allege that they did their job, why weren't the other functional groups able to take the ball and run with it. What if there wasn't a market for the feature or product in the first place, perhaps because the company never investigated their *marketability* before authorizing the R&D investment? Can the alleged innovation still be considered such, in any meaningful way?

One simple approach to understanding what's going on under the cover of "R&D" and to begin to measure its inputs and outputs more accurately is to temporarily set aside the magical but abused word "innovation," and instead analyze each "new thing" in terms of five possible classifications and definitions, as part of the quest to assess an individual company's *Return on Innovation*:

1. **Differentiation** = Innovation projects that yield clear, sustainable, and valuable differentiation against all known competitors, as evidenced by *proven* market acceptance and, more importantly, customers' willingness to pay premium prices. (Proof of market acceptance is an interesting topic for discussion in a separate article.) In essence, companies have become past masters at gaining *phantom* adoption for their new offerings by adopting such tactics as bundling them with their more established offerings. They then play the numbers game to report counts of "new customers," many of whom agree to take the bundle but have no intention of installing the newer, unknown component, and haven't made any conscious investment in it either.
2. **Productivity** = Innovation projects (or parts thereof) that do not differentiate the company in any measurable or significant way, but enable customers to gain productivity in their use of the company's products. While this result for customers justifies some innovation investment by the vendor, it is only marginally rewarding because customers are generally unwilling to pay a premium price for something that they see as a cheaper, better, faster benefit that competing vendors are likely to be able to offer sooner or later. The danger to avoid with this type of innovation is over-investment. Tech companies find it hard to do this due to their tendency to focus heavily on product features and functionality, especially when a competitor has beaten them to the draw.
3. **Neutralization** = Innovation projects (or parts thereof) that do not differentiate the company nor yield greater productivity to users, but enable the company to cancel out an existing disadvantage against its competitors, either for lack of a feature or because their feature performs suboptimally. A good example might be a kludgy user interface that is beginning to cause lost or slower sales against competing offers that provide better UIs. While justified as a relatively small proportion of a company's development dollars, this mainly comes under the title of 'sustaining engineering' rather than innovation engineering. Not surprisingly, this area can be a notorious sinkhole for so-called "innovation" resources.
4. **Failed attempts** = Bona fide innovation projects that, when originally conceived, targeted a real innovation opportunity, but took too long to bring to market and lost their differentiation potential to competing vendors, or for some other reason failed to produce any significant results in terms of productivity, or neutralization, per the above definitions of each. In a high performing organization, a reasonable percentage of failed attempts should actually be a sign of a healthy innovation culture that pushes the envelope. However, this privilege is often abused by underperforming R&D organizations with too much unmanaged bubble-up activity.
5. **Waste** = Innovation projects that resulted in wastage of resources due to being ill-conceived or poorly managed. This category also includes those projects that failed to

go far enough to achieve their original objective. When this objective was to create defensible separation against all known alternatives, the cost in terms of wasted resources can be astronomic.

Based on a wide sampling of systems and software companies, our firm has estimated in approximate terms that truly differentiating innovations usually account for less than 20% of the entire innovation budget, and sometimes even as little as 0%; at the other end of the spectrum, anywhere from 30% to 60% of innovation expenditures result in either failed attempts or pure waste. The difference between the latter two classifications, in my view, is that at minimum failed attempts should – if they are appropriately managed – result in positive outcomes such as important learnings for the next similar project, whereas wasted efforts generally leave teams and management demoralized, in addition to R&D budgets being drastically cut, people losing their jobs, and so on.

Back to the first and most valuable criterion, differentiation. Because effective innovation is not a game of financial or other numbers, we don't believe that it's reasonable to expect any company to truly innovate – in the sense of 'differentiate' – all of the time.

Furthermore, investments in productivity and neutralization, which are often bunched under the header of "sustaining" innovation, are absolutely necessary and should be expected to absorb their fair share of any company's overall R&D budget, particularly if the company has a full portfolio of offerings that includes some fairly mature ones. That said, the key thing for management to avoid is confusing differentiation – which should receive privileged level and quality of resource allocation – with any of the other four buckets.

So when it comes down to it, is it misguided to do the fashionable, democratic thing, and let innovative ideas, products and services bubble up through the organization, or is it wiser to return to the former times of innovation being more tightly managed in a top-down manner? My answer is that it is infinitely preferable to stimulate innovation from every feasible corner of the company. What is not appropriate, however, is (a) to define and describe it in terms that are too permissive of misunderstanding, (b) not to manage it with congruent metrics, and (c) above all, not to make the tough choices early enough in the development or launch process to recognize whether you are heading for an intended or unintended positive result (i.e., any one of the top three – differentiation, productivity, or neutralization), or a negative one – in which case it's important to pull the plug immediately in order to refocus resources and brainpower on new possibilities.

(*) See *"Dealing with Darwin – How Great Companies Innovate at Every Phase of Their Evolution,"* by Geoffrey Moore (2005, Portfolio - Penguin Group)

2. Case Study: Why it is Vital to Understand Customer Resistance to New Innovations

Three stories about technologies that are at different stages of adoption caught my attention in today's paper (SF Chronicle, 07.24.06). They each teach us distinct lessons about the dynamics that either accelerate or delay customers' buying decisions. One was about voice activation software adoption, the second was about the growing tussle between email and the IM-text messaging duo especially among younger consumers, and the third was about increased efficiency from managing household appliances remotely via home computer networks (the main example of actual use cited in the article was about a an apparent "domestic happiness-critical" application – nothing less than remotely monitoring the weekly laundry wash while performing other activities). I shall leave the last two stories aside for the moment, and focus here on the lessons to be learned from the first one.

Adoption of Voice Activation Software

The writer of the syndicated article, Brian Bergstein, uses a creative idea to begin his story: the first paragraph turns out to be what the software program wrote out after he dictated his message

to the computer, and we discover that the program committed no more than one apparently simple error – it left out the indefinite article “a” before the word ‘computer’. However, the lack of the “a” did alter the meaning of the sentence sufficiently to cause a double take.

Bergstein goes on to enumerate a number of advantages that the software (Dragon NaturallySpeaking 9, from Nuance Communications) brings to the table, and you get the impression that he is about to declare it ready for any user to buy and install. Despite these favorable factors, he goes on to dispute the vendor’s claim that this latest version of its software, just released, will make anyone who uses it happier and more productive. Bergstein suggests that the software is likely to be welcomed by the customary consumers of this type of program – people with disabilities, sore wrists, or voluminous typing responsibilities.

But, he points out, there are a number of disadvantages that add up to Dragon not yet being sufficiently compelling for the average user to opt for abandoning their tiresome keyboard just yet. These include apparently minor issues such as having to enunciate with an “inhuman level of precision” to avoid the program making its share of mistakes, having to spell out corrections carefully, getting a slightly delayed response when scrolling through or entering text in web pages, and the odd freezes and crashes of the program.

My diagnosis is that, despite what is probably the company’s hope that, finally, they have managed to come out with the most advanced voice-activated functionality available (which would no doubt be an admirable feat of engineering), they may still have had the tech company’s tendency to think *too* rationally about the benefits they offer, and too scantily about the “slight” disadvantages that could prevent users from making a buying decision. Customers need a single compelling reason to buy that outweighs the real and imagined disadvantages that almost always accompany a new-ish technology. And let’s remember, voice-activated software is still not widely accepted by end users as a de facto item of their daily lives, exactly because it is not free of inconvenient glitches and friction points. If anything, from a consumer’s standpoint, kludgy and unfriendly voice-based systems have been foisted on them mainly by service providers such as banks, credit card companies, airlines, and phone companies, who are anxious to cut their customer service costs.

Lessons Learned

It’s never fair to generalize – except when at least you recognize that it’s never fair to generalize! – but allow me to make the following observation: Vendors still get so mentally tangled up in the complications and difficulties of getting their products to the next level of acceptability that *they consistently fail to invest the same amount of thought to really understand the array of choices that customers have to deal with before deciding to adopt any new tool in their already crowded day-to-day lives.*

Another way of looking at what vendors often see as end user “inertia” (more often than not in a disapproving tone that suggest that it’s the user’s fault for just not getting it) is to remind oneself of the choices that customers have vis-à-vis how they respond to a company’s offerings:

1. *“Buy it now”* which requires (a) an actual decision to buy, (b) (normally) some kind of financial outlay or commitment, and (c) some kind of commitment to change what they are doing today, based on a strong emotional conviction that the product or service being offered will help them be more successful or productive than they can be if they miss this opportunity;
2. *“Consider buying it later”* which requires only thought about the three factors above, but no immediate commitment or action – and is therefore one of the “easy” options;
3. *“Don’t buy it (or anything) now”* which requires a simple decision to not change anything or make new any financial outlay or commitment for the time being – this is often the easiest option;

4. *"Don't buy it now or ever!"* which only requires a decision to banish your offering from their array of possibilities;
5. *"Buy something else"* that helps them solve their problem in a different or better way, which requires them to focus their decision and energy elsewhere;
6. *"Buy something completely different"* that solves a completely different problem and brings them different benefits, which requires them to spend their money elsewhere but still not on your offering.

Think about this list of options for a moment: Only one of the six choices results in good news *now* for the vendor; all the other options range from a vague "maybe sometime in the future" to "no way, no how, not ever!" When a vendor actually takes their marketing hat off and puts their "now I'm a consumer" hat back on, they can of course – at least, with a bit of effort – experience their own similar responses to being marketed at by sundry vendors on a daily basis. But it's a long way from appreciating what it feels like when you are a target of other companies' marketing or sales bombardments, to remaining sensitive to customer perceptions and feelings when the shoe is on the other foot and you are the vendor.

I am convinced that, if more vendors were to recognize the imperative to make their offerings so compelling *in at least one dimension* that customers are left with no emotional or rational option but to make choice #1 above, they would multiply their "hit" rate by a factor of 10x or more. One of today's best examples of a successful, and highly creative thought process around how to take a technology that was previously marketed poorly (MP3 players) and mold it into a thoroughly compelling offering – so much so that it has already achieved iconic status in the tech industry – is the Apple iPod/iTunes offering. And, make no mistake, Apple's success in this regard is as valid for enterprise-focused vendors as it is for consumer companies, because the need to empathize with the customer's needs and choices is important in every marketing or selling dialogue.

3. Dealing with The "Black Box" Syndrome and "Meeting Hell"

In this issue we are calling attention to a couple of massive sources of inefficiency and waste in today's tech organizations, particularly in publicly traded companies most of which are wilting under the increasing demands of serving their customers, giving their investors what they want, and fending off scrutiny by government regulators. If only they can find ways to become more transparent and more efficient in the way they conduct their daily activities, they will be more accurately evaluated by the market (which, come to think of it, may or may not be considered a desirable thing by certain companies!).

Challenge #1: Solving the "Black Box" Syndrome

For decades there has been a pernicious culture afoot in tech, and in many companies it doesn't seem to have changed much over the years. I call it the culture of the 'Black Box.' It is highly prevalent and extremely malignant, especially in its tendency to hinder a company's growth, conceal massive inefficiencies, and drain an organization's energy and resources.

So, what exactly do we mean by "Black Box"? What defines a black box for me is a lack of transparency and accountability, and – to cap it all – an absolute lack of shame about it.

Disturbingly, it is omnipresent in the two functional areas that in most tech organizations are by far the most empowered, R&D and sales, which only helps to exacerbate the problem. When the head of sales and/or the head of R&D tells you that you don't really need to know how they do what they do, because they are smart and they know what they are doing without having to follow a formally defined process (a dirty word in tech, since it forces orderly and predictable behavior in order to achieve desired results for the company), what is a CEO to do, especially if they manage in a relatively hands-off way? Part of the problem is that if the wizards in sales and R&D were to explain what they do, their ability to produce magic would vanish along with it.

But it doesn't stop with R&D and sales. I have seen many a company in which virtually every function is unmanaged like a black box; not only R&D and sales, but also professional services, marketing, finance, and operations. Not surprisingly, the more black boxes you have in any organization, the riskier the situation for the company. Why? For one, when problems hit, there's little or no way of knowing what's wrong, or how to fix it.

Among the unmistakable symptoms that your organization may be suffering an acute case of *Black Box Syndrome* can be found in responses produced whenever someone from outside the Box wishes to inquire as to progress in achieving their financial or other goals, or wishes to understand why something important did not get executed on. See how a typical senior practitioner of this somber art deflects all arguments for formalizing what their department does:

"What do you mean, do we use a repeatable project management process (or sales pipeline management, or systems development)! We don't need a process *per se*. We have a bunch of very experienced, smart people, who do highly creative work on a daily basis. There's no precise way of tracking or measuring it, precisely because it varies so much. In any case, if I spent my time micromanaging my guys around here, half of the smartest ones would up and leave."

Thus every request for information about how things get done in the organization in question provokes a "how dare you?" type of response, or – at best – a bluff to the effect that "what we do is pretty complex" (i.e., too rarefied in its cleverness and insight to be explicable to mere mortals like you), and it varies according to the customer or project (this is a common evasion).

There are ways of infiltrating these boxes and solving the puzzle of just how they get their work done, and how can I find out why they fail and why they succeed. However, it entails some fairly pragmatic activities, which I'll leave for another article.

Challenge #2: Dealing with "Meeting Hell"

The bane of life in every tech company, large or small, these days seems to be the endless stream of meetings, be they conference calls or actual in-person sessions, at which their presence is required. To make things worse, because every meeting is timed to last either 30 minutes, an hour, or multiple hours(*), no allowance is made for ending them late, nor for the need to move conference rooms, or buildings, or locations. Consequently, late starts are inevitable much of the time, and before long people come to expect late starts and late endings, at which time sloppiness often becomes the rule. Not only is this disrespectful and irritating to those who take the trouble to arrive on time and ready to start, but it generates inefficient and demoralized behaviors throughout the organization.

Most of our clients, whether they be start-ups, midsized, or very large companies, experience the same syndrome. It's a tech industry thing. And it suboptimizes every company's activities to an amazing degree. I would estimate that companies lose as much as 40% productivity via the following sinkholes: (a) poorly defined meeting objectives, (b) unclear accountability for running the meeting and/or for executing outcomes, (c) lack of a clear, feasible agenda or target outcomes, (d) poor preparation and, even worse, weak commitment and follow-up to agreed action items, (e) too many attendees who aren't sure why they need to be in the meeting (f) consequently, too many people semi-engaged or unengaged participants, (g) too many electronic gadgets switched on (for people to overtly or surreptitiously catch up on email or other more interesting activities).

So why aren't companies doing anything to solve the problem? And, in any case, what is the solution. Here are a few suggested remedies:

1. Instead of having people attend meetings on a need-to-know or not-to-be excluded basis, only invite those people who will need to do something substantive as a direct result of the meeting.
2. Arrange meetings for 25 minutes or less, 50 minutes or less, and so on, but never for whole or half hours. Allow time for summarizing the notes from the meeting, for taking a breather, bio breaks, walking from room to room (or building to building), and so on.
3. Make it clear who is responsible for running the meeting and, if appropriate, what kind of follow-up activities you are going to ask participants to commit to.
4. For larger-group gatherings, provide a brief, clear agenda with stated objectives, *at least half a day* before the meeting or call; request that each participant read the agenda and send you any additional topics they would like to cover *before* the session. Incidentally, it is still appropriate for longer sessions, such as all-day offsites, to spend some time at the beginning of the meeting to refine the agenda further, but you still want to make sure that everyone knows exactly why they are there in the first place, what they intend to get out of the meeting, and what they'll be expected to get done as a result.
5. Send any presentation materials at least two hours before the meeting, preferably earlier, to allow people to look through them beforehand.
6. Make sure electronic doodads are switched OFF during the session – or, in the case of Blackberrys, that people don't do the BlackBerry Hunch in your session – on pain of the perpetrators having to defer to decisions reached the actively-engaged participants. In return, if appropriate allow sufficient break times for people to catch up with urgent messages and activities.

At the end of the day, it shouldn't be necessary to adopt all of these measures in order to conduct business with a marked improvement in productivity. On the contrary, I would suggest that any two of these steps could make an enormous positive difference, without any need for a repressive or adversarial stand off between meeting coordinator and participants.

() It's a weak excuse to blame the technology in this case, but it does seem as if the calendar in the ubiquitous Microsoft Outlook/Exchange system is partly responsible for the slavish way in which meetings are scheduled and timed.*

Under the Buzz offers commentary on enterprise software and systems business and management issues. The goal is to provide provocative and accurate insights into the latest events and thinking shaping this continually evolving technology sector. *Under the Buzz* also provides commentary on strategies for building sustainable competitive differentiation and maximizing market valuation.

©2006 - Philip Lay

Disclosure: From time to time, the author and/or his firm may hold investments in, or provide advisory services to, one or more companies cited in this newsletter.