C3B Consulting's

INSIGHT-SPECTRA

Applying common sense to using technology intelligently

www.insight-spectra.com

Volume 25 Report 2; April 2012

Snatch and grab iPhone theft provokes rethink about smart devices and their eco-system lock-ins

Charles Brett, C3B Consulting and Constellation Research Group

Management introduction

On a pleasant spring day in Tel Aviv I was surprised to have a moto with 2 young men aboard drive up behind me and grab my earphone/mic cable that was attached to an iPhone 4. They then raced off down the street at speed, disappearing with their snatched treasures.

This iPhone's theft and deciding what its replacement should be, however, has started a process of reflection about the new smart device lock-ins that come from eco-systems like iTunes, the Android Market/Google Play, etc. Initial conclusions are uncomfortable for those who lose, have stolen or merely wish to change their smart device OS. Furthermore the replacement looks increasingly like being a stupid-smartphone rather than a super-smartphone.

iPhone considerations

Despite the iPhone/iCloud's self-locating capabilities, the stolen iPhone 4 is not likely to reappear (if it does that will be entirely to the credit of the Tel Aviv Police, who were courteous and helpful). After the usual processes of checking backups, changing passwords, etc., an uncomfortable awareness dawned — the simplest replacement solution would be another iPhone (albeit the expensive choice). Why? With a new iPhone one should be able easily to replicate the old iPhone environment via iTunes onto a new iPhone — apps, data, everything (assuming that iTunes works as Apple promises — and you only discover whether this is true the hard way).

But is a new iPhone still appropriate? The iPhone 4S does not really add anything special in capabilities (sorry, Siri) over the iPhone 4. Whereas the previous iPhone 3G had been an eye opener for its time (nearly 4 years ago) and swiftly became indispensable as a general purpose tool, the arrival of tablets changed the emphasis. Tablets are so much more usable that the need for a super-smart phone diminishes (which is a profound justification why something like the combination ASUS Padfone — including smartphone, tablet and keyboard — may yet become a big hit when it eventually arrives).

The opportunity cost of switching platform

Yet switching to another smart device OS has a cost — the app and content foregone cost. This is where mobile device eco-systems are delivering a new lock-in (in much the same away that IBM, HP, Microsoft, Oracle, SAP and others have ensured so well over the past 50 years for enterprise software).

Apps and content — whether for Android, iOS or Windows Phone — may cost a fraction of enterprise software but their investment adds up. I was genuinely surprised to find that my investment in purchased iOS apps and content for the iPhone over several years, when added up, was roughly as much the device itself. Thus there is a powerful financial incentive to remain with the OS you own, with its associated eco-system.

The difficulty is that though many apps have equivalents on iOS and Android and increasingly on Windows Phone, you need to purchase them again if you change OS. This lock-in contrasts with Amazon: its Kindle Reader is available on multiple platforms (Android, iOS, OS X, Windows, Linux, etc) and your purchased from Amazon content (though not apps) can then move around irrespective of platform. Such flexibility has real benefits, and even more so in a mobile enterprise environment.

Early initial candidates are, therefore, for either for a Padfone (because of the combination of capabilities mentioned above, and even though it has Android) or a Windows Phone (for reasons that will follow). Android, in principle however, is less desirable because of the number of OS bifurcations (2.2, 2.3, 3.1, 4.0, etc) — plus too many Android vendors do not provide OS upgrades for their devices.

Yet another reason not to buy another iPhone

Though unwanted, the re-think about what should replace the iPhone has continued, with new and unexpected directions of thought. While the easiest solution is still another iPhone, that remains the least

likely option. Instead, the concept of a 'stupid-smartphone' is coming into focus and is taking hold (not that such an oddity yet seems to exist).

I still have my original iPhone, and initially thought of using that. But, as it is of the 3G generation, it cannot run iOS 5. This is a major problem. Indeed I am not even sure what would happen if I did connect the iPhone 3G to the iTunes I used for the now gone iPhone 4 — but I am 99% clear that something would have to give, with choices being dictated and implemented by Apple rather than me. In other words the case for no new iPhone continues to grow, even if it means I will take both financial and practical hits when losing the apps that I had previously run on the iPhone 4.

An iPad changed the requirements

In addition a more fundamental reason for not re-signing up for an iPhone has emerged. When I first bought that now elderly 3G it was mind-opening in what it could do. Being able to communicate (email and talk and sms) and research (using the Safari browser) as well as run apps when on the move was a revelation. Moving up to an iPhone 4 made perfect sense.

But then the iPad arrived. I waited, until buying one seemed to make sense. Almost immediately my habits changed. The iPhone was no longer my first choice for checking email, researching/browsing, reading and running apps. Within days (not even weeks) the iPad had supplanted the iPhone as the mobile device of choice, except of course for phone calls and for the data plan that came with the iPhone. In effect the iPhone became just a phone with additional sophisticated capabilities that I could use if I needed them, but this was less and less often. The acquisition of both an Android and a Windows tablet (both touch screen) confirmed this impression: the iPhone 4 had evolved into an expensive device where only a minimum of its capabilities were frequently used — of which the principle ones were general communication-enablement (voice, sms and tethering) and size (as in, it fit in a pocket).

Now I am without that iPhone I am not as bereft as I would have been 2 years ago. After a fortnight of studying what is available my initial inclination is to go for a Windows Phone. I already like Windows 8 on the Fujitsu Q550 tablet as well as Windows 8 Consumer Preview (Win8PC) on a PC. If Microsoft can deliver what it discussed at the Win8CP launch at Mobile World Congress (MWC) in Barcelona in February — namely the ability to sync across all of phone, tablet and PC-with-Win8 — this will be a major benefit. If apps on a Windows Phone can run as apps on a Win8 tablet and even (like gadgets or widgets) on full Win8, that would be a real bonus. (Of course whether this will work as described will remain an unknown for some months, or longer ...)

In addition, and this applies to Android as much, nearly all the functionality in those now-lost iPhone apps is available on Windows Phone. It may not be the identical app but there are few apps that really matter which are unique to the iPhone platform. Yes, paying again is a painful necessity. But platform portability across Windows-based hardware vendors has a major attraction (some software lock-in seems always to be inevitable). Yet my analysis has moved on still further ...

Roll on the stupid-smartphone; forget the super-smartphone

Now back to the stupid-smartphone. What I am now broadly concluding is that I do not need a top-rated, super-smartphone any longer.

What I do need is a phone which has some smartphone characteristics, but not everything that an iPhone or an ICS Android device or even what a full Windows Phone 7.5 device offers. Thus far my list of critical if limited requirements (and none of these are state of the art) includes:

- making voice calls and sending sms messages (for the device to be a phone)
- the ability to share a data plan — tethering/acting as a router for both WiFi and 3G or 4G connections so that tablets, PCs and other devices can exploit one phone-company data plan and without each device having to have its own data contract (this concept should have worked a treat for me at MWC, except that Movistar (Telefonica) and Orange Spain were not competent enough to provide data services at the mobile world's premier event — see http://bit.ly/yot7jY)
- a decent browser, which means reasonable response/rendering speed, size and quality screen presentation (but nothing as good as Retina or the latest AMOLED alternatives)

- some limited, specific applications (including email, contact management, etc)
- a camera, but not a high quality or sophisticated one (more for reference than taking artistic photos)
- some form of entertainment capability (for music, radio and podcasts but probably not for video or TV, where tablets excel)
- limited storage; best of all would be some form of high speed microSD that would enable data (and even some apps) to move between machines.

Pretty much everything else that I may want will run or be available on a tablet (iOS, Android or Windows Phone) or PC (whether Win 7 or 8 or OS X or Linux). In effect the processing power and storage needed in a smartphone reduces once you have a tablet. Even better, the limited capabilities of a stupid-smartphone (by comparison with what is available today in most smartphones) would turn out to be significantly cheaper (and be less attractive to thieves in Tel Aviv, or elsewhere) as well as satisfy my practical needs.

Management conclusion

The obvious immediate candidate for the stupid-smartphone is that original iPhone 3G. Sadly the inflexibilities of Apple (and its iTunes) render this option a non-starter (plus the battery, which is not user replaceable, is no longer up to it after almost 4 years).

Early 2012 was probably not the time to 'lose' a smartphone. The mobile device world is in a state of flux:

- *Android and iOS, unless you root or jail-break, bind you to Google and Apple*
- *the Windows Phone/Windows 8 (for all of smartphones, tablets and PCs) remains imminent even if Windows 8 on a Fujitsu Q550 tablet looks great*
- *there are signs of Linux coming to smart devices (tablets and phones)*
- *other ecosystem players are around (of which Amazon is only the biggest one so far).*

On the other hand, mid-2012 may be a good time to decide what will be the replacement phone. The choice looks like being between a Windows Phone or an Android one, but which is still unclear. Fortunately I do not think I need to make a decision for 2-3 months. But I am increasingly sure a lesser smart device (that stupid-smartphone) which can act as my mobile comms. hub is the form of solution that will be most appropriate. Indeed, I suspect that the same will apply to most people in enterprises — once they have started working with tablets.

Charles Brett
C3B Consulting and
Constellation Research

What is a 'Data Scientist'? And who is not one?

Neil Raden, Constellation Research Group

Management introduction

There is a perception among larger organizations over the past five years that more quantitative methods, with or without Big Data, are critical to business success. The problem is that most commercial organizations have little to no depth in analytical disciplines.

In this analysis, which first appeared as a Constellation Research blog posting, Neil Raden considers:

- *what is a Data Scientist*
- *who is one, and who is not*
- *how should they be trained*.

The term 'Data Scientist' over-reaches

In those businesses where data and data products represent the primary revenue stream there is an abundance of talent. Some, like Google or Amazon, employ hundreds of applied mathematicians and statisticians, in the same way that manufacturing companies employ mechanical or electrical engineers. Medical informatics, genomics and even intelligence and defense groups work on the bleeding edge of research into methods for classification, prediction and optimization. This work is often unique, involving massive data volumes accompanied by unruly data formats and sources that are beyond the typical enterprise's data flows. Coupled with a broader understanding of the business or organization, a name for these professionals has emerged — 'Data Scientist'. This seems to me to be inaccurate.

Let us look at how this plays out. The work currently undertaken is clearly divided between that of:

- true scientists, those who research and create algorithms and methods, publish papers and actively participate in their discipline's communications
- those who understand and employ quantitative methods, design, test and deploy models — but do not create new science.

I refer to these two as Type I and Type II respectively (in a research Report to be published later in 2012 by Constellation Research, I will go into much more detail and also introduce describe Types III and IV). In my analysis the first group (Type Is) are truly the scientists. In contrast the Type IIs are not, even though this is the group typically referred to as 'Data Scientists'.

Type Is

In practice there are very few true Data Scientists in commercial organizations. Data Scientists work in research, academia and organizations where the production of new methods and algorithms are the core of the enterprise.

Google, Amazon, Wall Street and a small number of others are exceptions. These are examples of businesses where Data Scientists do indeed produce new methods in quantitative science as well as publish in peer-reviewed journals. But these are the minority (of enterprises).

Type IIs

Although, there is a prejudice that tends to be in favor of employing PhDs as Type IIs, this is not a necessary qualification. the reality is that, despite the unfortunate name given to this growing class of professionals, these people are not truly Data Scientists.

But that does not mean what they do is without value. Indeed their efforts represent a new sort of vital data analytics role in organizations. Yet finding suitable people is difficult because the role needs all of the following capabilities:

- familiarity with varied types of data, and the resultant multitude of analyses employable
- skill sets which not only include programming but also use of quantitative methods, investigative techniques and comfort with modeling
- an ability to discern between what is meaningful and what is not
- sufficient domain knowledge (in contrast to being a quant-for-hire)

- comfort when communicating complex subjects to others (who may lack the background in the tools and methods employed).

Not all is theory

I mentioned engineers briefly. Qualified engineers often start work with a solid theoretical grounding in the area of their choice, but often start work with no real practical experience of engineering — and typically no experience at all in the business domain of their employer. They learn as they progress. In fact, there is even a professional designation for engineers which demonstrates they have the skill, training and practitioner's experience to be a senior engineer — Professional Engineer (usually abbreviated to PE).

A different model for recruiting and nurturing professionals for the data analytics role, instead of competing for a small pool of PhDs who may be overqualified and unfulfilled with the work, might come from their way insurance companies grow their own actuaries (Mr. Raden is familiar with this because, he has an actuarial background). Indeed there are two major actuarial organizations in the US:

- the Society of Actuaries
- the Casualty Actuarial Society.

Both organizations administer comprehensive (in practice, grueling) certification programs that start with most of an undergraduate math degree and proceed to all aspects of probability and statistics plus understanding the insurance business itself. The exams can take 5-10 years to complete, and most insurance companies offer time off for study as well as on-the-job mentoring.

Two aspects about this approach are key to their success:

- first, a Fellow in either Society has to demonstrate a thorough grounding in quantitative methods
- second, and perhaps even more importantly, you only become a Fellow of either Society once you have also demonstrated a true understanding of the workings of both the business dimension as well as of the entire insurance industry.

Analytics is growing in importance and deserves something similar.

Tons of gimmicky professional 'certifications' exist today (often associated with specific products rather than general principles). In contrast an actuarial fellowship, Professional Engineer certification or even a CPA (for accountants) all represent the completion of rigorous, practitioner-oriented programs. Enterprises cannot, and should not, expect universities to provide this kind of education.

At the same time it is manifestly true today that skill with data, and analytics, has already become central to much if not all of organization's success. It is time to be serious about analytics.

You can call these people 'Data Scientists' — but businesses will have to participate in their learning and training. Whatever you call them, they do not grow on trees. Some sort of legitimate professional certification is needed. But until then, enterprises need to take grooming and nurturing these professionals seriously.

Management conclusion

So what do we call them? Certified Data Scientist? I do not think so. A Fellow of the Quantative Analysts Society? Better.

Data Alchemist? I like it, but it is not really comprehensive enough.

If you have any other suggestions, let Mr. Raden know (you can follow him on Twitter — @neilraden).

Neil Raden
Constellation Research
www.constellationrg.com

Reflections on BYOD (Bring Your own Device) — Part II

Larry Fulton, Enterprise Architect

Management introduction

Larry Fulton has worked in IT for 20+ years as an enterprise architect, both as a user and as a provider of knowledge worker tools. This gives him a practical perspective which spans both sides of the Bring (or Buy) Your Own Device (BYOD) debate.

In this, the second of two analyses (the first appeared in the January, 2012 **INSIGHT-SPECTRA***), he reflects on BYOD and considers how IT will have to respond.*

How can IT be expected to provide support for, well, *anything*?

With BYOD one immediate concern is the resource cost of providing support for anything that a user happens to want to use. Within the traditional computing environment, IT already supports numerous configurations, numerous versions, and so forth. Though IT probably limits the end-user environment to one operating system and two or three hardware vendors, the possible combinations of hardware, software, infrastructure and applications are many. As new devices and applications continue to emerge, the challenge for IT continues to grow.

The issue that IT fears is that it will be expected to help make almost anything work in the corporate environment, whether or not IT has seen it before. The underlying question is: where should (or can) IT draw the line? With a fixed support budget, there has to be some point where IT is able to say: "we just can't help you with that". Right? Perhaps.

Where is the enterprise's data? Is IP protected?

Moving on, there are some truly challenging questions that require enterprise consideration before BYOD arrives. Employee-owned devices create the certainty that company data, once upon a time safely ensconced within the corporate technology environment and behind corporate boundaries, will or could be stored in hundreds or even thousands of devices over which the enterprise may have limited control. To complicate matters, employees do not always understand what information belongs to the enterprise and what information does not. When everything was stored on enterprise disk drives they did not need to consider the issue — and any confusion had few consequences.

In the case of intellectual property, as opposed to simply company records, the issue of visibility and control becomes even more complex. External copies of internal documents add more challenges.

Even more pernicious, consider any instance where enterprise people are working on project deliverables while using their own devices or web-based tools to create these. The resultant documents may never have existed on enterprise resources (like storage drives)? How can companies protect their legitimate interests in this environment? How can employees protect theirs? To whom does the resultant IP belong?

Finally, how can one protect an enterprise's information assets from misuse, theft or disclosure without some means to ensure its security? This is a high-stakes game.

It is difficult enough when you know where all the data is kept. When it is potentially copied or replicated to many BYOD devices that an enterprise may not even control, the problems become more significant. Furthermore, that is even before considering regulation and compliance requirements.

Do not forget another unknown — clouds

A list of IT concerns would not be complete without mentioning the many variations of 'cloud-based services'. Knowledge workers love services like EverNote, DropBox, Box , iCloud or 37signals, to name just five — because they make it easy to get things done. The number and variety of services that include a cloud-based component is exploding. Smart devices are part of this expansion — because people want to be able to look at their 'work' (whether personal or enterprise) on smartphone and/or tablet and/or laptop. Take a look at EverNote: it is easy to share files

between home and work environments? DropBox accounts are free. Need a backup solution because neither you nor your IT organization has anything in place to protect you against the failure of your company-provided laptop or your personal Android tablet? Multiple online, cloud options are found in minutes.

All of these solutions share several key attributes:

- they store your (or your company's) information on cloud-based servers that neither the employee nor the employer control
- most of the cloud service providers have legitimate needs for their software to read data to fulfil the needs of the service they provide
- too many do not encrypt the information they store (and, even if they do, they often have the means to decrypt it if necessary)
- while it is generally in the service providers best interest to take great pains to prevent security problems, that does not mean they do not happen anyway.

Even though IT also cannot guarantee that the enterprise data it controls is 100% protected, that is not the same as entrusting it to a third party where there is no legal relationship with the enterprise itself. If an employee 'contracts' with a free cloud service, there is no exchange of value contract with either the employee or employer organization.

This has led, understandably, to a widespread and reasonable concern about anything that moves or stores data in anything that might be called 'the cloud'. As a result, many enterprises are choosing to block access to known cloud-based services from within their corporate networks. But this will not work outside those networks, like when an employee connects from home or on the road.

So what is to be done? Establish new types of relationship?

These trends are not especially new. Many talk about the need for IT to:

- be proactive
- work with users
- anticipate developments as well as problems
- find ways to give users what they need before they go out to find it for themselves.

The simple truth, however, is that IT will never successfully put itself at the front of all these issues. The market forces unleashed by BYOD and mobility are going to continue to create new, attractive capabilities that users will wish to harness. Knowledge workers will simply go ahead and take advantage, because they find real business value and because they can buy their own devices and support services.

Others suggest that IT should be working with suppliers to ensure that IT's security concerns are being met. The Windows platform has catered to IT in this regard over the years, and to some extent device and platform vendors will make changes to address the most significant problems.

But the BYOD/mobile world of today is a new world where business users make their own personal acquisition decisions (and write the checks). IT's influence will be strongest in those areas where IT as an industry speaks with a common voice or has an investment that BYOD users need — for example for corporate Exchange/SharePoint compatibility and/or associated policy enforcement.

In a sense, the world in which IT operates is not changing; it has already changed. In effect there is a small and narrowing window of time during which IT has some power to limit what is happening in its environment. IT needs to use that time to think differently about BYOD and mobility — and plan for the future we are going to receive, instead of the future IT may think it wants.

Look down the road, perhaps as long as three to five years out, but maybe even sooner than that. Knowledge workers will acquire the tools they find the most attractive. If your organization makes that too difficult, they will go elsewhere.

The real question for enterprises concerns how will they adapt to such a reality; will:

- they enable their workforces to be more effective?
- key elements of workforces drive increasing risk of data exposure?

IT needs to re-think

IT must rethink how to fit knowledge workers into corporate technology environment. Increasingly tools will be hosted in environments outside of the control of corporate IT, whether on employee-owned devices

or in clouds. Many devices will be new and unfamiliar.

Remember, however, this is not that much different than the world IT supports today for an enterprise's customers. IT already provides web-based applications, in both traditional and mobile device flavors. IT provides mobile applications that work online with corporate servers to provide richer customer experiences. Many even provide public APIs for the direct use of customers to support third party innovation.

Rather than looking at knowledge worker tools as something that must be supported directly within the corporate environment, IT should instead look to support knowledge workers as they connect to the corporate environment. This has another analogy. Airlines used to think they were in the business of flying planes. This led to major customer disenchantment, until airlines realized they were in the business of transporting people (which had an impact for a while). IT needs to re-think, to focus around supporting people to do their jobs rather than running IT. Put another way, IT cannot control what tools will be used; but it can control the connections and interfaces employees use to work with enterprise resources.

In this context, IT needs to become real about information sensitivity. IT cannot afford to default everything to being treated with the utmost security. By establishing clear guidelines (policies) and communicating them to knowledge workers, IT can achieve a higher level of protection for the data that is truly sensitive. This is where IT needs to work to build an alliance with knowledge workers, and work with them at the policy level, so they understand what should and should not be done. IT cannot afford to understand every device and tool that will be used. IT has to elevate the discussion to the risk and policy level, where employees take responsibility.

CIOs must stop IT people trying to establish meaningless or unworkable policies. At the same time, while the general principle should be to promote a spirit of cooperation with users, enterprises need to be clear where they draw the line. If inhibiting the copying data to the cloud is a strict policy, then it needs to be communicated in detail — with second offenses, at the least, dealt with severely. (Too often, policies are so poorly worded and poorly explained that this ends up forcing an enterprise to be lenient of well-meaning interpretations.)

The emphasis must be on taking responsibility, with everyone understanding where enterprise-ownership boundaries really are. Once understood, this encourages innovation within the chosen boundaries. When this occurs, BYOD can contribute and provide clear business value.

Management conclusion

As Mr. Fulton argues, knowledge workers have a responsibility. If they wish to leverage new tools and technologies, they must make themselves aware of the implications for their employers. Equally, employers must facilitate this.

It is easy to expect IT to be responsible for protecting each enterprise's information assets, but it is unrealistic to expect IT to do this alone. IT simply cannot keep ahead of all of the developments that knowledge workers will encounter and adopt.

Simultaneously, IT has its own a responsibility to:

- *avoid unnecessarily limiting what knowledge workers can do*
- *find ways to help knowledge workers continue to innovate on behalf of the company (and a big part of that will be finding new and effective ways to educate knowledge workers and enlist them in the task of protecting each enterprise's information assets).*

The most successful enterprises in coming years will be those that attract the most forward-thinking knowledge workers, whether within IT and elsewhere. This means encouraging a style of work that knowledge workers find rewarding and productive. In effect this means BYOD is here to stay.

Larry Fulton
Enterprise Architect

Coping with traffic surges

Peter Bye, Bye Associates

Management introduction

In an earlier analysis on planning for disasters ('Is there an ostrich in the house?' **INSIGHT-SPECTRA** *, Volume 23 #3, May 2010) Peter Bye observed that many organizations fail to make adequate, or even any, provision for recovery following a disaster. Given that the potential consequences for any such organization could be catastrophic, the continuing failure to think ahead surprised then, and continues to do so.*

Since Mr. Bye's original piece, there have been some spectacular disasters: hurricanes, earthquakes and, by far the worst, the tsunami in Japan. These events not only take a terrible human toll but they also wipe out data centers, thereby crippling IT services if adequate disaster recovery plans are not in place.

But it is not just 'natural' disasters that compromise IT service delivery. Sudden traffic surges can cause equally dramatic increases in the workload on the systems used, often bringing them to their knees. All the components and their supporting infrastructure are intact — it is just that they cannot cope with the demands placed on them. In this complementary analysis, Mr. Bye looks at the effects of traffic surges and how to mitigate them. He begins by examining example causes.

Common causes of a traffic surges

Most IT systems experience significant but expected fluctuations in load, typically associated with dates or events. Month- and year-end processing or big increases in travel associated with holiday seasons are examples. For these there is time to plan — although too many systems still seem to manage to crash even when the extra load is anticipated long in advance.

An instance of the latter is the annual clearing of entrance applications to UK Universities. The key school examination results are announced in August. Students who do not obtain the grades needed to go to their chosen university have the option to use the University clearing system to look for alternative places and courses. The demand is hardly unexpected. Even so, last year the system failed.

It is, however, the big increases in traffic where there is little or no advanced warning that cause the real problems. Natural disasters are often the trigger. Incidents such as floods or earthquakes put tremendous loads on the IT systems used in the response. Emergency services are a prime instance. Police, fire, ambulance and even military systems will likely be involved — depending on the scale of the incident. In practice all these are likely to experience and have to deal with processing loads that are far greater than their normal traffic volumes.

Man-made overloads also have their impacts

The same is true if disaster events are man-made rather than natural. The riots in England during August 2011 put substantial extra loads on the emergency services' systems. The Olympics or World Cup, though scheduled years in advance, offer additional examples of short duration, man-induced peak loads where commercial and other public sector systems may also be affected.

In all of these, telecommunications as well as similar infrastructure industries are likely to see major increases in traffic volume when an 'event' occurs. Following Hurricane Katrina, the number of faults reported to the telephone company covering the area increased dramatically. Recording them and subsequently scheduling repairs sharply raised the workload on the systems concerned. The same applied to the Loma Prieta earthquake, when the telephone system (more fixed then than the mobility of today) was overwhelmed by people calling out of the quake area, as well as people calling out, to see if relatives and friends were OK.

Less cataclysmic events than hurricanes and earthquakes can also cause sudden surges. Special promotions for products or services, discounted fares by airlines for instance, increase the demands on systems. Such promotions are often announced at very short notice, to surprise competitors. The business making the promotion should have been able to prepare its own systems (assuming that marketing and IT

did not underestimate the resulting demand) but other organizations affected — such as for payment processing — may not have had the same warning.

The Internet and the steady increase in mobile device usage (such as with smartphones) is now playing its own role in creating extra traffic. The load on a system obviously depends on the number of people generating it. In the case of a discounted fares promotion before the Internet customers would have contacted the airline or a travel agent. The relatively limited number of people in the airline or travel agencies able to make bookings for customers would have had a throttling effect on the volume (albeit at the expense of frustrating customers trying to make contact). In contrast, the number of self-service devices (such as smartphones or tablets or PCs) able to access today's systems is, if not infinite, becoming very large indeed and able to generate instant huge traffic volumes.

As the world becomes increasingly dependent on IT services, the disruptive effects of natural or man-made disasters are becoming more severe. The problem of unexpected traffic surges is not, therefore, going to disappear. Rather it will worsen.

Planning for the unplanned

When thinking about the unplanned it is essential to recognize that events causing traffic surges are not unexpected. It is when they will occur that is unknown, or at least hard to predict.

Therefore, just as plans should be made for disasters, provision should be made for sudden, unplanned workload increases. The next question is: what form should such provision take?

The fundamental requirement is for a degree of technology support to exist which is elastic enough to cope with major variations in load. One possibility is to provide systems powerful enough to handle just about any possible peak. But cost becomes a problem. Users are then paying for a big enough system in which the capacity available would be under-used (hopefully) for most of the time. In general this is not financially nor operationally acceptable.

The next question: is it possible to provide sufficient capacity for peaks while keeping costs to a minimum? One approach is shown in Figure 1.

In this approach, systems are commonly implemented in three functional tiers handling the user interface, application and database respectively — a process that may be called vertical distribution. Each tier can also comprise one or more servers or virtual servers — which may be called horizontal distribution. The grey shaded boxes in Figure 1 illustrate a normal configuration. Additional power for peaks is added by

Figure 1: Adding power by adding servers or virtual servers

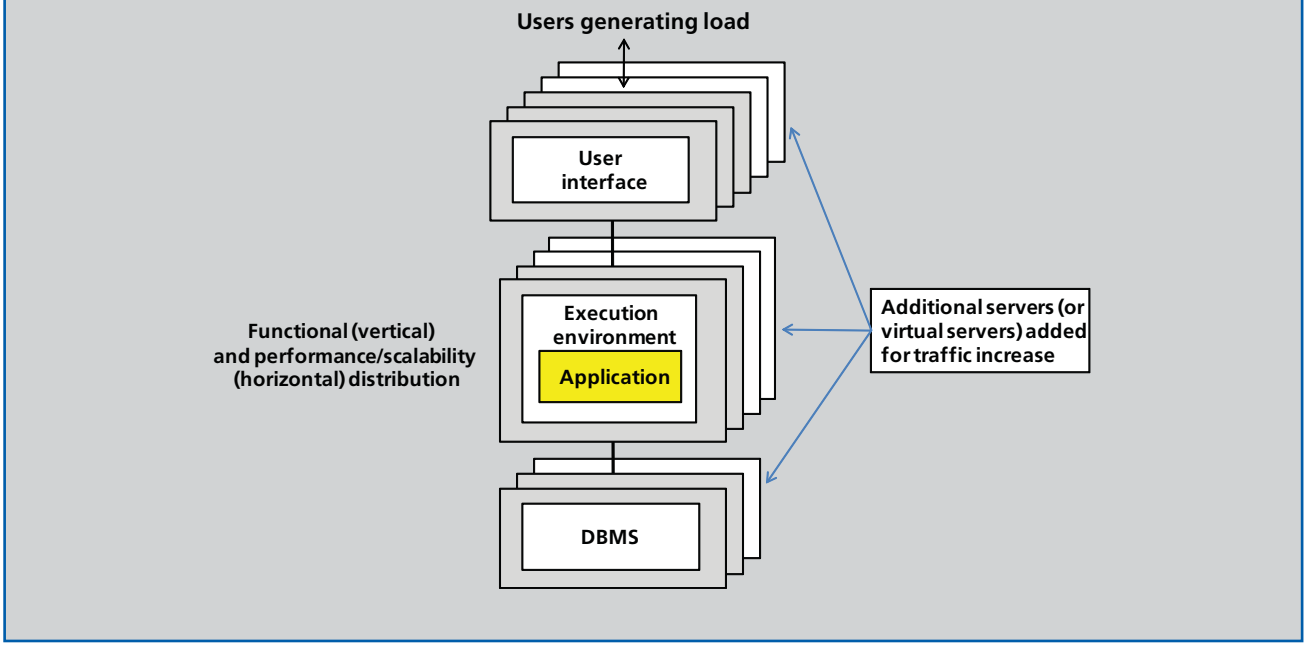

switching in more servers or virtual servers, as shown by the white boxes.

This process can become overly complex, requiring a high level of preparation and then automation. Loads must be monitored and, on reaching a level where overload is likely, the needed extra servers or virtual servers have to be taken out of a pool (or from other, less important, applications) and repurposed to support the overloaded application.

A third approach, which may be unique, is that introduced by Unisys with its ClearPath systems. This is essentially a 'pay-for-use approach', with metering technology. Technology developments, especially multi-core chips, have now reduced hardware costs, thereby enabling the economic delivery of systems with far more power (MIPS) than is needed for normal operations while still at relatively little investment cost (for either vendor or customer).

Excess hardware leaves headroom available for planned peaks or sudden shocks. The customer only pays for what is used (in 'MIPS time' units) — not the actual full capacity available in the system. (This is analogous to electricity supply. The incoming supply allows wide variations in load. Customers pay for the Kw hours consumed.)

In Figure 2:

- the grey-shaded columns show the normal usage over a day, for example, with peaks and troughs with the average consumption, and the power required to handle normal peaks, indicated (an unmetered system would require that level of power to be able to handle normal peaks)
- the white columns represent the extra load caused by an unplanned surge.

Although this surge could occur at any time at any level Figure 2 illustrates the worst possible situation — when the surge load is added on top of the normal peak. An unmetered system would require this level be marked to handle exceptional loads. Customers with ClearPath systems can instead configure systems so that they have a maximum possible power of several multiples of the normal or average — and much more than normal peaks. Thus unexpected surges are swiftly absorbed, without human intervention.

As long as surges do not develop into an extended extra load (say several weeks — or more) the long-term average need not be significantly affected. Overall it is the average the user pays for, not the occasional peak.

Management conclusion

As with making provision for IT (or other) disaster

Figure 2: Pay-for-use, using metering

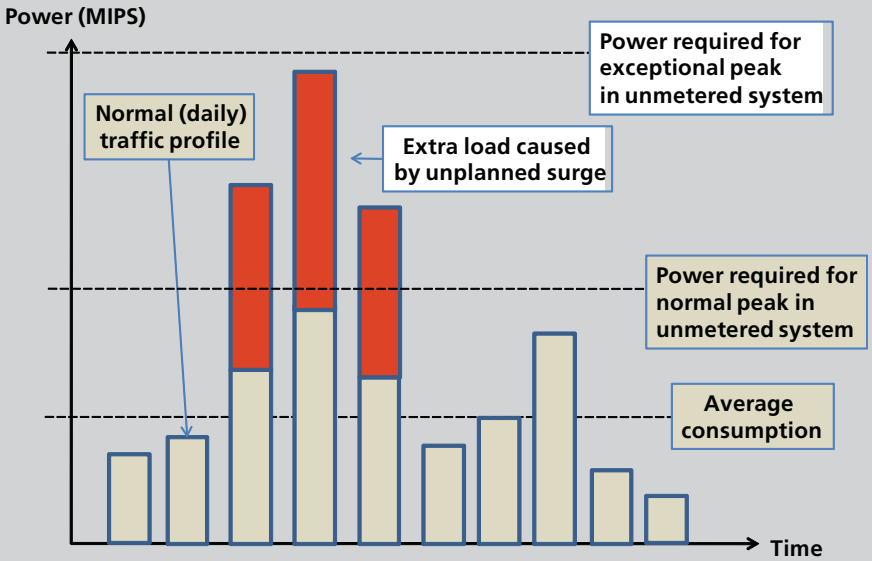

recovery, it is essential to plan for unexpected traffic surges resulting in serious system overloads. Such provisions should be commensurate with the degree of exposure.

Critical systems — such as those used in emergency services — must be capable of responding immediately. Indeed, extra unexpected loads may follow from a serious external problem such as an earthquake or typhoon or hurricane — making it even more critical to respond effectively. The last thing needed in an emergency is for response times to extend, or system instability to occur, caused by overloading.

Of the three approaches, Mr Bye favors metering as the simplest and most effective. It immediately handles sudden increases in load, because systems are designed to absorb this. There is no need for intervention, repurposing or any other activity to provide the needed extra capacity. There is, therefore, less risk of some additional error occurring in the process of adding extra processing power.

Metering also brings a beneficial side effect. The extra processing power available in the system reduces the elapsed time required by many processes such as the execution of a transaction. The power applied in MIPS is higher but the elapsed time is reduced. The power consumed in 'MIPS time' units is the same as for a lower powered system but the work completes faster.

Peter Bye
Bye Associates

IT transformation's next step: mobility encourages enterprise cloud platform adoption

Robert B. Cohen, Cohen Communications

Management introduction

Mobile devices appear increasingly to be battering down the traditional walls of corporate IT systems and in so doing encouraging enterprises to create new designs and architectures for computing and services. When vendors dictated (as was usual in the past) which systems enterprises used, there was little need for interoperability. The arrival of masses of intelligent mobile devices has shattered this presumption.

Not only are these mobile devices forcing change on previous practices. They are also opening the opportunity for new visions of enterprise IT for the future. In this analysis, Bob Cohen looks at what is changing, and the implications.

A changing environment

In some firms the decision has been made to (try to) limit the use of smart devices to IT staff and/or top executives only. In most cases, such a policy is more about trying to figure out how the enterprise IT environment should respond and then manage them (rather than inhibit use). In contrast other firms are moving to transform their overall approach to IT, using mobility as the game changer.

What is at the bottom of this? Is it the management of diverse mobile devices? No, it is a future where the vision is that enterprises — as well as consumers — will want everything delivered 'as-a-service'. The old IT model was just not able to evolve in this way. Most likely, single vendor solutions will also not be the way

to respond to the challenge that mobility and mobile devices present.

Services via cloud platforms

Providing services via cloud platforms means some cloud platforms will be inside the enterprise. But even important internal private production cloud platforms for the enterprise can be hosted and managed by service providers (rather than IT) from outside.

Other platforms can connect into a single distributed cloud and are not related to production activities. Public external clouds have the additional advantage that they can respond to consumer thinking and comments, even if they are also operated by service providers.

But the net of either approach is that functions which IT had previously thought its own preserve will largely be outsourced in the future. Indeed, increasingly service providers look more and more likely to shape and even determine what technologies are chosen for supporting cloud platforms.

One immediate downstream effect is that the role of IT in purchasing for enterprises will decline. Another is that telcos may yet play a far more dominant position in determining how to architect cloud platforms, develop cloud services and support both production and non-production environments for the enterprise than IT had anticipated.

For enterprises, this will produce a very different approach to architecture. Many firms are already in the middle of migrating data off traditional approaches, like Oracle, and shifting it to open source data storage. In effect they are assuming that the timing is already right to begin a significant transformation and to create new platforms that are not available from commercial vendors.

The mobile dimension and service provision evolution

In mobile, one of the emerging groups that has not had much impact to date is the WC3 Core Mobile Web Platform Community Group, though it has developed a new mobile browser test suite — Rigmark — that was released in late March. This is an attempt by more than 30 device manufacturers, carriers and developers to accelerate the standardization of mobile browsers. This will let developers see what browser supports the functionality that an app needs. The Community Group has also been working with service providers on a way to streamline billing.

What these efforts do not address, however, is how mobile devices access enterprise apps and data. With standardization it is possible that enterprises might be able to force vendors to adopt new browser requirements. But equally it is not clear that these mobile efforts will dovetail readily with the aims of future-oriented cloud platform proponents in the enterprise.

That said, the cloud platform for the mobile enterprise will possess an architecture which supports a service management-based model, plus the use of a services catalog. IT managers within firms, consequently, will more and more need to focus on managing how cloud platforms function and deliver (rather than doing the implementation themselves). They will be concerned about:

- the distributed cloud management layer
- the interfaces between different service provider cloud platforms and enterprise applications and services.

Assuming this vision for cloud platform model is successful, service providers will need quickly to become skilled at operating hosted private cloud platforms as well as private external cloud platforms and future external public cloud platforms. This will present challenges for service providers. They will need to move rapidly:

- from hosting and applications provisioning
- to running cloud platforms in collaboration with cloud service vendors.

Some of the relevant vendors will be 'older players', like Alcatel-Lucent and Huawei. Others may be firms that telcos have not significantly partnered with before, such as Salesforce.com. In addition, service providers will need to create new business models to evaluate continuously the profitability of their cloud platforms and the outsourcing of cloud services development and support. This will mean understanding the mobile and applications focus of recent years along with more of a focus on cloud applications for business customers, both large and small.

Is change happening?

Is there evidence that businesses plan to move to

clouds in a more open way and to do it more quickly than many might assume? Some businesses such as Wal-Mart have set up a special group (@Walmart Labs) to consolidate new technology efforts and move the focus from ecommerce to efforts to change the way the company uses IT. Similarly, the Pistoia Alliance of pharmaceutical firms has released a series of requests for vendors to build cloud computing projects that can be exploited by four major firms.

These are, as yet, relatively small initiatives. But they are likely to be indicative of what will happen.

Management conclusion

Mr. Cohen believes that the challenge for service providers is to shift their current focus to cloud platform management and distributed cloud management for enterprise customers. This will almost certainly require a transformation of existing skills and business practices of service providers, which is likely to be a continuing process over the next ten years.

Nevertheless, service providers increasingly seem likely to be the main drivers of changes in the enterprise (and towards IT) in the future. This represents a very different role to that they have played in the past.

Robert B. Cohen
Cohen Communications

INSIGHT-SPECTRA

is published and distributed worldwide by:

C3B Consulting Ltd.
19 St. Michael's Road
Winchester SO23 9JE
UK

Telephone:
+44 787 233 4000
+34 686 116 993

Email:
insight-spectra@
insight-spectra.com
or
charles.brett
@c3bconsulting.com

World Wide Web:
www.insight-spectra.com