

C3B Consulting's

INSIGHT-SPECTRA

Applying common sense to using technology intelligently

www.insight-spectra.com

2

Message to Iberia (and others): IT to customer communication needs to fly First Class

Charles Brett, C3B Consulting

4

Operations and infrastructure professionals: still worried about your job? You should be.

Larry Fulton, Consultant

7

IBM's OPAL: open treasure or hidden riches?

Charles Brett, C3B Consulting

8

Reliable application/transaction processing in clouds

Joe Clabby, Clabby Analytics

11

Value-based case study #4

A private cloud pays back for a financial service provider: 25% fewer servers, yet 3x more workload

14

Value-based case study #5

Bluhalo: exploiting the cloud to cut costs plus hugely increase capacity and resilience

17

Value-based case study #6

D&B innovates by integrating coAction's external SaaS platform with its existing internal applications to create a new customer Collections product

Volume 23 Report 2; March 2010

Message to Iberia: IT to customer communication needs to fly First Class

Charles Brett, C3B Consulting

Management introduction

Information, information, information. This, when used intelligently with technology, is the key to limiting damages and even improving customer satisfaction in a crisis.

Seems obvious — no? Not apparently to Iberia, several other airlines and, C3BC suspects, many other organizations which assume a crisis will never hit them. Yet the solution is so simple and the benefits are so great that it beggars belief that large organizations in 2010 can be both so insensitive and inept as well as customer unfriendly.

A bad experience opens eyes

Let me begin by saying that the experience (described below) does not qualify as anything special and is trivial when compared to the miseries in Chile, Haiti and elsewhere. Flight delays and execrable airline service describe the lot of most modern, frequent air travelers (and many have seen far worse than what happened to me). But this experience set me to thinking about minimizing the impact of a crisis — on suppliers, staff and customers.

In January a combination of bad weather and an airline ill-prepared for systemic problems (and apparently uninterested in its customers) meant that a journey with one connection that was scheduled to take 9 hours unnecessarily took 33 hours — with much waste of time and effort. This was only the second time I have traveled with an airline whose internal disorganization under difficult but not unusual circumstances virtually incapacitated it (it had happened to me once before while flying Air Canada, which I now avoid when I can).

This time it was Iberia at its headquarters airport in Madrid. As is so often the case, it was not just snow (the reported reason) but multiple compounding factors, including that stunning lack of interest in, or understanding of, its passengers (and, C3BC suspects, its staff) that made everything so bad.

To understand why a solution is so straightforward,

let me first offer some background. At 0230 on the scheduled day of travel, I received an SMS message from Iberia saying my flight from Alicante-Madrid (at 0715) was rescheduled (but with no details) and that I should call Iberia (with no number given to call). Now awake, I tried the Web to find that Iberia's web site said all was working as scheduled. Without a number to call, I had no choice but to leave for the airport as if all was working. On trying to check in at 0600, I was told that not only was my flight cancelled but so was the previous one — because the planes had been unable to leave from Madrid the evening before. In other words, Iberia knew it had problems many hours before it sent that SMS to me.

More distressingly, I discovered that no attempt had been made to rebook me or the 200+ other people now queuing before 2-3 stressed Iberia reservation staff. Eventually, I was rebooked for an 1150 flight to arrive in Madrid at 1300 with my new connection scheduled nearly 12 hours later at 2330 that night.

I checked in and went through security when my 'troubles' really began. There were NO Iberia employees beyond security for passengers to consult. That meant relying on airport (not Iberia) TV monitors or the good will of a lounge. What information existed was almost always inaccurate. Even when I turned to the Internet, I found that Iberia's own Arrivals and Departures pages lagged other travel sites (like Flightstats.com).

Other contradictions abounded. For example, one kindly Iberia manager who called from my ultimate destination to update me, said he had just been told I would not be flying at all that day; yet I was on a plane 30 minutes later (and even that was not full, despite three or more flights having been canceled).

Eventually, I reached Madrid (barely 250 miles /400Kms away) in the late afternoon (some 10 hours behind schedule) to find Terminal 4 (Iberia's new prestige home) in similar confusion. My misfortunes continued when all passengers had to disembark from the 2330 flight at 2350 into a deserted terminal (again to an absence of information) — because the

plane would not work. Finally I arrived at 1600 the next day — merely 24 hours later than scheduled.

Information absence and the consequences

What I wanted most — and I suspect most Iberia staff wanted — was accurate, timely and updated information, preferably with some context (like some ongoing explanations). C3BC presumes that what an organization wants when facing such dire straights, is the time to solve its operational issues. In essence, it wants breathing space to address and correct its difficulties.

Without information, paying customers (passengers) become a cost rather than an asset. In wanting/seeking to know more, they lose sympathy and proceed to further stress the supplier (the airline) with continuous questions. Furthermore, those employees who are in the crossfire — in this case facing irate and frustrated passengers — often have no or only incorrect information, which only compounds the confusion (they are asking the same questions, on behalf of customers).

The situation with Iberia was a recipe for trouble, which is exactly what Iberia duly served up. Yet the keys to preventing and correcting such situations are not novel. Information plus communication buys that breathing space that the supplier needs. It does not take much wisdom to spot that if you deliver constantly updated information, you keep customers informed, which relieves the burden on customer-facing staff members, who in turn do not spend their time trying to obtain the impossible from the hard-pressed operations people trying to resolve the problems. With a constant flow of updated information, especially for an overtly customer-unfriendly airline like Iberia, the odds are that customer (and even staff) relations will markedly improve.

The irony is that implementing a solution for such circumstances is relatively trivial. As C3BC research and analysis indicates, all that is needed are: 1) a dedicated Web site, 2) a browser, 3) SMS and 4) a specific 'crisis' information communications group.

By dedicating a Web site to the 'crisis' an organization:

- removes substantial query traffic from an organization's main order/sales Web site (for

Iberia, its reservations system), which organizations need to keep running in order to accept ongoing (future) business

- establishes some sense of responsibility for those affected by the crisis, along with the impression that someone is doing something, that the situation is evolving and improving and that the organization cares (as British Airways has done in the past)
- exploits access to real-time information by making it almost universal — via a browser, whether on laptop, computer or phone, either at the airport or further away.

Of course, having a Web site is no good by itself. This has to have as much updated information as is practical. In Iberia's case, a small 'crisis' group should have been tasked with keeping the crisis Web site up to date (this includes liaising with operations swiftly and simply). This group should include the technologists who can rapidly process relevant information onto the Web site. Personally, I would add a commentator to give a running update about the current position (improving the sense that events are moving forward). A more superior Web site might even assist customers to do their own rescheduling online — after all, they probably know best what they need. All of which, if implemented, reduces the burden on other hard-pressed resources and staff.

The third technology is SMS (in this I include Blackberries, instant messaging as well as social networking sites like Twitter, Ping! and other near-immediate communication media). Recall that SMS received at 0230. Iberia would have helped much more if it said the flight was canceled due to bad weather (we all understand this happens), that I was rebooked on the XYZ later flight to connect with the 2330 flight and/or that I should call a given number if this was not acceptable. Had I been sent this information, I could have stayed at home until late afternoon (though it would not have helped when the plane at 2330 would not function).

As an aside about SMS, while Iberia's performance was execrable, it is certainly not alone. About 18 months ago, I was talking with the IT integration staff within the operations group at a major US airline. The idea of alerting airline colleagues (never mind customers) by SMS was greeted with puzzlement — even though these able people know of operational problems (like an aircraft would be leaving late) before anyone else. The notion of using operational

data to alert passengers (never mind colleagues) was regarded as abnormal.

Of course, not all affected passengers (and people waiting to meet them) have the necessary technologies to access such alerts or the Web site. But what I have noted time and again is how fellow customers share information generously. If you receive an SMS telling you that your flight (or order) has been rescheduled yet again, you will share this with others around you. In this way (what might be called real-life social networking), customers help each other, but they can only do this if they are given the means to do so. Conversely, customers (usually inadvertently) will hinder if left ignorant — as Iberia discovered.

Management conclusion

Looking back from the outside at what happened to me, Iberia made just about every mistake in the book. Had it done even a little prior planning, once it knew it had problems the previous evening, it could have easily forestalled the avalanche of bitter complaints

that I overheard. Regardless of how bad the news is, passengers (customers) prefer being informed to being left in the dark.

The net of this is that:

- *most organizations, with a little forethought and planning, can exploit simple technologies to make themselves work better (and look better) in a crisis situation when all appears to be going wrong.*
- *the alternative — failing to take simple steps — can prove very expensive, not just in lost customer satisfaction but in squandering client, as well as, staff loyalty (as Iberia is discovering on its bottom line).*

Charles Brett
C3B Consulting
www.c3bconsulting.com

Operations and infrastructure professionals: still worried about your job? You should be.

Larry Fulton, Consultant

Management introduction

Although the global economy continues to show signs of modest recovery, IT people who have kept or found jobs would be wise to hold off on any celebratory sigh of relief. Past layoffs and increasing pressure to perform have made many IT jobs increasingly stressful. Unfortunately, the continuing pressure to drive costs down is not going to abate any time soon.

While it is probably true that many organizations have now moved beyond a reflexive cutting of jobs to contain costs, many companies are now taking a more methodical look at where they can drive costs out of their operations through restructuring, optimization, outsourcing and other means. In this analysis, Larry Fulton considers various implications and their ramifications.

Flipping the 80/20 cost picture — a priority for everyone in IT

Corporations want to operate as leanly as possible across the board, but IT has its own place in the crosshairs for several reasons, including:

- renewed interest in changing the legendary 70/30 (or even 80/20) split between operating and maintenance costs and the delivering of new capabilities; one interesting survey from Information Week (December 21, 2009) listed 'flipping' of this ratio as a top 10 priority for CIOs in 2010
- new CIOs (and other executives) needing to demonstrate their commitment to holding costs down; IT's big cost center' reputation is less and less popular with businesses every year, and CIOs with the ambition to

- play a larger role in the business are looking for ways to shrink apparently runaway costs
- the desire of businesses to take advantage of the slower economy — when times are slow and competition is less severe, many companies see an opportunity to move in new directions in order to be in a stronger position when business circumstances eventually turn around; these companies want to maintain or even increase spending on development, because new development is the pipeline for adding new business capabilities and pursuing those new opportunities.

The net: pressure to reduce costs in IT's operations, support, and maintenance functions will be ongoing.

Insufficient justifications and unfortunate truths

There are many reasons that lie behind the costs of a typical large (and small) IT operation. IT is complex. The evolution of most IT environments has been unstructured, at best. Supporting a complex environment built through multiple mergers and acquisitions, built upon countless generations of technical infrastructure and 'assembled' with multiple incompatible development approaches, is not easy and is costly. The number of moving parts in a modern business application is much larger than it was only ten years ago. The number of applications in large IT shops easily reaches into the hundreds or even thousands.

The unfortunate truth is that justifying these costs is no longer sufficient. IT infrastructure professionals must understand that executives now compare financial performance to industry benchmarks. Similarly, customers are evaluating operational performance against other organizations and technology press information. This is true whether you are a CIO, manager, team leader or individual contributor. If any part of an organization, or indeed the organization as a whole, fails to deliver, expect significant pressure to reduce costs while simultaneously improving services.

Prescriptions

It is straightforward to see that any prescription for success must include the following steps:

- doing whatever is necessary in the short term to deliver state-of-the-art service at a competitive price (I will come back to

- addressing the short-term challenges)
- articulating a long-term vision that continues to drive costs down and value up
- making the capture and improvement of metrics part of everything you do
- ensuring that accurate metrics are available to stakeholders (and that you regularly communicate them across the organization).

For a long-term vision, I recommend (from personal experience) looking at what is happening across your segment of the IT industry. Remember that the same cost pressures you may feel often spark innovation.

Take just one interesting example — storage (among the least glamorous of IT cost centers). Specialty vendors have introduced new solutions at significantly lower prices than traditional ones. This has caused the latter to offer new classes of products that sport more attractive price points, driving down the cost of storage technology as a whole. With vendors competing to offer new products at better prices, there are golden opportunities to reduce long-term costs and modernize at the same time. (see also the BluHalo case study on page 14).

Capturing and using metrics is, quite simply, a 'just do it' proposition. It is important, therefore, that you include both financial metrics — the cost and value of services you provide — along with functional metrics.

Communicating value to stakeholders is another no-choice scenario. What happens if you really are the low-cost, high-quality, forward-thinking provider, but that is not clear to your company's management? The answer should be obvious — no one can act on information they do not have. If decision makers do not recognize your facts you should probably assume they (the decision makers) are not receiving them.

Job one — perform at the top of the practice, now

So what can you do to shore up your short-term value proposition? When your management or your customers assess your organization's performance, the criteria are brutally simple: what do your services cost relative to other similar organizations/departments/teams, and what level of service are you providing relative to other organizations/departments/teams?

This means making comparisons to your peers within

the company and in other companies. The good news is this means you should be able to determine easily where you stand.

Whatever service you or your team provides, you need to know what it costs and how much it costs other organizations to provide the same service. Even if you do not have the answers, someone in your organization either has them or will obtain them. The real concern is how your costs stack up against similar organizations, using metrics like the number of servers per service administrator, the total support cost per terabyte of storage or the total cost of ownership of a Windows desktop. If your costs are higher, the first and only question will be: 'what are we doing to get these costs under control?'

Keep in mind that you may not be aware that such an investigation is under way — the people in the room may be limited to your executives, your accounting group (providing the only data that will be considered about the actual cost and the reasons behind it), and maybe even third party providers claiming they can beat your costs. Or maybe your development teams are being asked to improve their performance, and among their obstacles is their inability to have development and testing environments provisioned in less than three days. 'Why can't our teams deliver this level of service?' is the question that will be asked.

You probably have a long list of justifications for the service levels you do provide and the associated costs. These are usually a mix of actual limitations established by the very management teams and customers they serve and the inertia of doing things as they have always been done. It is probably true that no one cares: if other organizations can do better, why can't yours?

If your team is in this situation, you need to help overcome two fundamental problems: the team itself is almost always wrong:

- about which problems are fundamental and which problems are primarily inertia-related
- investments in its own seemingly time-tested processes and procedures.

Things that are 'impossible' often have a way of being done by someone else. Policies that limit your options may have outlived their usefulness; challenging them now may result in the ability to meet evolving needs more easily. You need to be aggressive in

finding out what customers really want, what they do not like and why you are doing things in a way that do not meet customer expectations. If customers want requests handled more quickly, find out why they are not handled more quickly today and what can be done to eliminate those challenges.

Processes established to manage customer requests are likely culprits — you may well see them as reining in the chaos. But customers see them as yet another set of hurdles controlled by long-lead-time support groups. Review procedures and talk to clients constantly. Team up with other support groups to consolidate these processes. Be responsive — processes should have stakeholders, not victims. Look at the services you are delivering: assess similar practices across your industry to make sure you are not falling behind the state of the art.

Management conclusion

Ongoing success in today's environment requires delivering state-of-the-art services at a competitive cost, evolving and then evangelizing a longer-term vision. This should take advantage of IT market opportunities and implement ongoing processes of measuring and reporting the business value you are delivering to stakeholders in your organization.

The bottom line is that operations and infrastructure services continue to be commoditized; there are many third party solutions and service offerings that will look like a better deal than the services you or your organization provide. In some cases, this is so true that businesses feel compelled to go for that better deal. In the cases where you are positioned to provide the best solution for your business at a competitive cost, do yourself and your organization a favor — see the writing on the wall, and start tuning up your operation so you are the best deal in town.

Finally, if you think something cannot be done, you need to find out how other organizations are able to do it — and then go ahead and do it. You have to focus on making your service better and cheaper than the best industry benchmarks, in the eyes of your management and customers. You cannot afford to allow anything to stand in your way.

Larry Fulton
Consultant, Analyst
and Enterprise Architect

IBM's OPAL: open treasure or hidden riches?

Charles Brett, C3B Consulting

Management introduction

IT software customers live in an increasingly expensive world. CIOs worry and worry still more as the escalating cost of software licenses and consequent maintenance consumes a greater and greater portion of their budgets. Towering above even these concerns, in C3B Consulting's experience, is the trauma and cost of integrating disparate pieces of software, particularly that associated with middleware, into operations management.

IBM offers customers one approach to this in the form of an extensive library of over 1850 — often free — 'integrations' that can be obtained over the Web. This library rejoices in the name OPAL (standing for Open Process Automation Library), which can be found at www.ibm.com/software/tivoli/opal

Treasure, or not?

It can be argued that OPAL is a hidden jewel which should shine especially brightly in tough economic times. Yet, on another level, OPAL is classic IBM: the jewels are hidden. Try finding what you want, and if you are not already Tivoli-educated (for OPAL is thus far primarily Tivoli-oriented), you may have real problems — even though the solution you need may well be there. The oddity here is that so much of what OPAL offers remains so hidden.

So what are these 'integrations' found in IBM's OPAL? From C3B's discussions, it appears they are pieces of code or applications which enable two or more pieces of software to work with each with each other, one of these being some part of the Tivoli family of products. Take for example: TelAlert for Tivoli NetView, a tool IT can use to deliver and communicate urgent information and thereby improve the management of operational work flow processes. TelAlert integrates with Tivoli's system management framework so that customers have global, secure guaranteed bi-directional notification to multiple devices — including PDAs, Blackberries, IVR and mobiles. Moving such information out to a variety of devices requires multiple protocols which usually means integration difficulties. Yet a solution is already available in TelAlert.

It is this immediate availability that attracts. Yet, what is more significant to C3BC is that these integrations more than just are the work of IBMers; ISVs and partners already provide over 20% of the integrations available. The real and potential benefits of OPAL offerings are, therefore, far-reaching.

According to IBM, integrations are used by end-user customers, as well as by ISVs/SI/business partners and even by the IBM sales force (in the latter case to identify solutions in advance to integration issues that will likely occur once software products have been installed). IBM also indicates that such pre-canned integrations are a significant factor in winning business — and C3BC can believe this: how rare it is to find that an integration solution exists before you encounter it in practice.

Analysis

Despite all the many positives, C3BC remains puzzled. The OPAL concept is broadly excellent but its public manifestation is narrower than one would expect. That ISVs (as well as IBM) have developed integrations clearly suggests recognition of the importance and value of such an (integration) library. Yet, at least thus far, even though these integrations are ready and available and low cost (or free, as many are), C3BC believes IBM seems intent on missing an opportunity. Indeed, it seems to C3BC that IBM should more clearly observe the literal meaning of OPAL — the Open Process Automation Library — by opening it to platforms beyond Tivoli. By doing so, OPAL could become a universal library of diverse solutions that are ready to enable integration.

In C3BC's analysis, this could be a door opener for Tivoli plus other IBM's Software Group products — and fits well with today's need to contain or reduce IT costs. With an already impressive number of integrations available, it is the integrations that should be the focus of OPAL: after all, simplifying integration issues is what IT users look for. With the increasing importance, and arrival, of the need to integrate technology to support a smarter planet, a greener earth, OPAL could become still more significant.

Management conclusion

In medieval times, “the opal was considered a stone that could provide great luck because it was believed to possess all the virtues of each gemstone whose color was represented in the color spectrum of the opal. It was also said to confer the power of invisibility if wrapped in a fresh bay leaf and held in the hand (Source: Wikipedia).” This strikes us as an excellent metaphor for the power and transparency that an open integrations library should achieve.

The outstanding question in our mind, remains, will

IBM capitalize on what it already possesses and take OPAL further and faster for the benefit of an even wider customer audience — and itself? Or will it keep the jewel of that is its Open Process Automation Library hidden from all but a limited audience?

Charles Brett
C3B Consulting
www.c3bconsulting.com

Reliable application/transaction processing in clouds

Joe Clabby, Clabby Analytics

Management introduction

Amidst all the cloud computing chatter of the past year or so, a few elemental questions have largely been ignored:

- *what exactly happens when a transaction is sent into a cloud for processing?*
- *which physical/virtual resources does it use?*
- *if you do not know which resources a given transaction is using, how do you approach capacity planning in a cloud?*
- *what dependencies does the transaction have?*
- *if the transaction is performing poorly, how can the fault be isolated?*
- *if the transaction misbehaves intermittently, how can that fault be isolated?*
- *how do you tune transactions in the cloud to improve performance?*

Until recently, Clabby Analytics would have argued that performing transaction processing within amorphous cloud environments was asking for trouble. It would have pointed out that the tools and methods used to track transactions across server tiers were archaic (network sniffer tools and time correlation methods). Plus, it would have also noted that performing cloud capacity planning and performance tuning is difficult if you cannot easily track applications across physical tiers.

Capabilities are improving. Clabby Analytics has changed its view. It now believes that a new generation of ‘Business Transaction Management’ (BTM) tools make it possible effectively to track transactions and rapidly troubleshoot problems within cloud environments. What has changed its analysis:

- *attending a Computer Management Group industry event in Dallas in December where several sessions dealt with transaction processing in the cloud*
- *speaking directly with practitioners who troubleshoot applications, perform capacity planning and tune cloud-residing servers and applications*
- *several briefings vendors to discuss their transaction handling in the cloud*

As such, this analysis provides a quick overview of the BTM marketplace and describes some of the challenges faced by information technology (IT) managers and administrators as they look to exploit cloud computing architecture. It also briefly examines one of the emerging leaders in the BTM field: OpTier.

Tracking applications in a cloud: a look back at Grid architecture

Clabby Analytics considers the distributed computing architecture known as ‘grid’ to be a predecessor to

cloud computing. Back in 2002-2005, we closely followed the grid computing marketplace and found grid solutions to be well suited for handling behind-the-firewall scientific and technical applications. They were also appropriate for addressing some business problems (provided said problems could be parallelized across many computers and that the conclusion of each parallel stream could be reassembled after processing to produce a final result).

But we also found several shortcomings with grids, related to the processing of certain types of workloads, including:

- complex batch applications involving a lot of data transfer
- heavily transaction-oriented applications (where transactions needed to be tracked and monitored—and potentially rolled-back in case of failure)
- applications where external databases needed to be accessed regularly
- application environments requiring strong authentication/authorization for security.

Because tracking down application/transaction failures was so complex a process, grid users gravitated toward predictable, parallel workloads such as those found in scientific environments. In these environments, if a segment of an application failed, it could be easily be restarted and reassigned to other computing resources. This capability has been proven by some highly creative ‘volunteer’ grid efforts, including SETI@Home (Search for Extra-Terrestrial Intelligence) and the cancer research project — Folding@Home.

Clouds and grids share much of the same basic architectural constructs — they are both computing architectures that rely on widely distributed, often virtualized resources to execute workloads. Logically, clouds should face some of the same challenges that grids do in terms of application management, performance tuning, capacity planning, etc. But, thanks to a new class of advanced BTM applications, clouds are becoming less prone to the same application tracking, monitoring, and management issues that had hobbled grids in the past.

What is BTM?

As is the case with many evolving markets, a concrete definition of BTM as a discipline is hard to come by, primarily because vendors tend to bend the definition

to fit their existing application and transaction management toolsets and strategies. Several IT analyst firms (including Gartner, EMA and IDC) and a few vendors, however, have attempted to provide working definitions for BTM (see Note 1 for the reference).

To Clabby Analytics, BTM is simply an offshoot of the application performance management (APM) market — a set of products and utilities that makes it easier to manage applications and transactions as they cross server tiers within a cloud. But one big difference between APM and BTM is that APM has been generally silo-ed within a single tier of the transaction flow (end-user, J2EE/.Net, messaging, CICS).

In contrast BTM ties together the silos of APM to make a full-enterprise view of the transaction. This union enables IT managers better to understand automated process flows. Accordingly, BTM tools are becoming increasingly important to enterprises looking for ways to build cloud environments. Not only do BTM tools simplify transaction management, they are also useful in helping to troubleshoot and tune business process flows.

What is more important than a definition, however, is what BTM enables IT managers and administrators to do. In short, BTM software can provide a graphic, topological view of an application as it travels across physical tiers within the cloud. In turn, this information enables IT personnel to observe application/transactional behavior as applications traverse cloud environments, interact with middleware and databases (this view is especially useful when it comes to troubleshooting). BTM can also provide a view of resource utilization that can be helpful in capacity planning.

BTM: different approaches

Many vendors claim to offer BTM software solutions that can troubleshoot, control, and manage transactions within cloud environments. However, only a few can actually deliver well-integrated BTM solutions.

The big difference in BTM vendor approaches can be found in how vendors find and isolate transaction flow problems. Some vendors measure activities in physical silos, using a network sniffer/time-based correlation; others use deterministic monitoring/management tools that can span server tiers while providing a topological view showing how each transaction is behaving in real-time, for each one.

The big problem with network sniffers is they cannot track a specific transaction traversing a network. They collect snippets of a transaction flow, which can be used to help isolate the location of a server (or a process on a server) that is experiencing a resource consumption spike or service elongation. But a sniffer cannot weave snippets into a clear, cohesive view of how a transaction behaves as it multi-hops through a morass of distributed servers; therefore discrete transaction problems that are low volume or that hang anywhere in a transaction life-cycle are missed. Extra time and work (and sometimes a team effort) are required to locate the source of these types of problems.

Other vendors' products (including HP's Transaction-Vision, MQSoftware's Q!Pasa, CA's Wily Introscope, OpTier's BTM Suite, and a variety of IBM's Tivoli products) offer varying levels of instance-level transaction monitoring support using a deterministic or content correlation approach. These can thread together and topologically represent an entire transaction as it flows through multiple systems in a cloud. Techniques vary in the level of automation and the level of effort that might be needed to assist in flow detection.

But these vendors do provide good insight into application/ transactional activities within a cloud. As a result, cloud managers and administrators can view an accurate representation of transactional dependencies, resource usage, service levels and other elements that pertain to transaction management. Armed with this data, managers and administrators can isolate problems in the application, database or middleware layers — and bring in the right human resources (from their application, database and infrastructure organizations) to fix them.

The OpTier approach

Of these vendors, we have seen the OpTier products demonstrated (we have only had briefings with several of the other vendors). What stands out to us about OpTier solutions is that they enable IT managers and administrators clearly to track cloud-based applications, and gather performance and resource utilization data. This simplifies tuning cloud performance and provides the data needed for capacity planning.

OpTier's approach is highly integrated and accurate — making it easier to use than comparable sniffer/time correlation methods. This integration, com-

bined with the accuracy of the data that OpTier collects, makes it possible to manage transactions effectively as they cross tiers in cloud environments. Furthermore, OpTier has been able to build a fully automated flow discovery environment, which may explain several statements made by its customers about delivering a positive RoI within weeks.

Management conclusion

Tracking application and transactional faults in distributed systems or cloud environments can be deviously complex — because this requires following programs as they hop from server to server making database calls, with reads and writes along the way. It also involves tracking business process flows to ensure that transactions execute efficiently. Fortunately, new BTM tools that support deterministic transaction monitoring/management are now available to help IT staff more efficiently and effectively manage their existing distributed computing environments and evolving cloud environments. Additionally, these tools also make it possible to prove clearly to stakeholders in application development, infrastructure and information systems organizations that a given problem is unambiguously theirs to fix.

When evaluating BTM tools (Note 2), Clabby Analytics recommends that enterprises ensure that the tools they evaluate are capable of monitoring/managing transactions across multiple tiers. This point is critical because troubleshooting transactional problems very often involves traversing many server tiers (and several BTM tools clearly do not have this cross-tier capability). Also, make certain the tool that your organization chooses is capable of capturing resource consumption at each tier a transaction traverses (this is important for cloud performance management and capacity planning).

Note 1: You can find the material on BTM definitions at:
<http://businesstransactionmanagement.blogspot.com/2008/10/definition-of-businesstransaction.html>

Note 2: For more information/reports about BTM and cloud transaction processing, see
[www..clabbyanalytics.com](http://www.clabbyanalytics.com)

Joe Clabby
President, Clabby Analytics
www.clabbyanalytics.com

Value-based case study #4

Private clouds payback for a financial services provider: 25% fewer server, yet 3x more workload

Management introduction

John Parkinson, who describes himself as a serial retiree, is Chairman at ParkWood Advisors LLC. Two years ago he was asked by a financial information provider to direct the strategy, setting of standards and architecture for a new computing production infrastructure. This resulted in the introduction of what today is called a 'private cloud'.

In the past two years this has produced, as he discusses below, significant results. This will only improve as future plans coalesce, to roll existing mainframes into the new cloud.

The business problem

The financial information provider where I have been working is an interesting place. It is all about information. Without information, there is no business. To collect, organize and then deliver that information, everything is technology-based. To achieve this we process consumer purchase data and then generate millions of reports and calculate billions of scores each day, for access by our customers. We organize masses of data into a big data warehouse and then we generate, each day, an operational data store against which we run applications to generate over 5 million credit reports and about 1 billion scores.

The basic problem I was handed was that the business needed to address the degree of variation between the base and peak loads of the computing environment required to run daily operations, which was seen as a rising expense that need to be controlled. Peak compute loads then ran at between 5-10 times the base load. We had large amounts, therefore, of much underutilized compute-capacity. Typically, for example, if we had 100 units of base load, our average usage was 250 units of load — with a peak of almost 1000 units. Thus we were sometimes using almost a 1000 units but far too often only using 100 units.

A further point to make here is that our average utilization, even after years of optimization and experience, a 15% utilization rate for much of the compute

capacity. By industry standards this was quite good for a non-mainframe environment. But it was not good enough and even with cheap compute cycles this represented an inefficient use of capital. The business, therefore, wanted a better return on its computing asset investment. My challenge was to architect and build a computing infrastructure that could rapidly (measured in tens of seconds) be reconfigured to cope with both planned — and unplanned — surges in compute demand.

In addition, we had a techno-cultural issue. Traditionally, our product teams — who built the applications that process the information which we sell to customers — created these applications in silos. Anytime anyone wanted a server, they tended to buy what was necessary. Applications, in other words, tended to have their own dedicated hardware. As a result we had many individually configured discrete servers — about 1100 of them from 3 different vendors. There were servers that were 10 years old and others that were 10 minutes old. We had 7 or 8 different flavors of operating system. Although each individual acquisition had been warranted and justified, in aggregate our approach was inefficient.

Where we started ...

We started by adopting 3 overtly simple principles — borrowed from Randy Mott's (HP's CIO) thinking about 'ruthless standardization' — that everything that can be standardized should be standardized. We adopted the following three principles:

- Principle 1: standardize what we buy
- Principle 2: virtualize everything possible
- Principle 3: automate where possible (we wanted to have the minimum of human intervention determining how computing work was being allocated and performed).

We made the decision to standardize on HP for server hardware. We then sat down our and HP's architects to work out an approach to meet our objectives.

Two years ago there was no reference architecture

that we could use for a challenge like this. We decided to proceed incrementally — starting with a relatively small number of (processor) cores and see where we reached.

As part of moving forward we decided to standardize on blades (one issue you discover when building a cloud is that there are an awful lot of wires involved, with connecting these being an error prone process). We then decided to apply Principle 2 and start by virtualizing as much of the essential connectivity as practical as well as the compute infrastructure. We acquired two chassis, each half filled with blades and implemented using HP's Virtual Connect technology, Sun's N1 Grid Engine, VMware's ESX for the virtualization and Red Hat Linux. We then built up our software stack on these — as the cloud underpinnings.

Once installed we started to throw workloads at this infrastructure — albeit in a test environment — with the deliberate intention of comparing production system costs with what was running in our new cloud. We were looking to compare what we were paying for virtualization and for the grid manager as workload overhead with what utilization we were obtaining in normal production plus we tested how fast we could reconfigure by adding or subtracting capacity.

What happened?

As almost always occurs when you are proceeding without a roadmap — and even knowing where we wanted to reach — some things worked really well, some things sort of worked and others did not work at all. For example, we found that the way we managed workloads through our cloud platform (the grid manager) worked really well. The automation side was, therefore, pretty good, although scheduling was still a challenge. In contrast, the failover and recovery was so-so: most of what we wanted worked but not always as described (this was usually fixed by going back and changing the documentation).

The part that was most difficult was the rapid scale-up/scale-down capability that we sought. After 2 years of experience we can now add cores into the operational mix pretty reliably on an on-demand basis — but only from a known, pre-allocated pool. We have found it much harder to add brand new capacity on the fly: that is not yet sufficiently reliable for us.

Today we have about 8000 cores available. If we add another rack of servers, which can be up to an addi-

tional 1000 cores, it still takes about 2 days before we can be sure that all those additional cores are being made available in a reliable and productive way. The issue is that the resource manager cannot react fast enough to the addition of that much extra capacity. Instead it goes into a constraint loop as it tries to rebalance the workloads. This is a non-trivial computer science problem. But we are confident a dependable solution will emerge in time that we can get steadily better at this.

Virtualizing is more than creating an image

Initially we started by virtualizing a copy of existing servers. This virtualized copy was then run in our cloud. In essence nothing about that server had changed except its 'execution home' which was now in the cloud. What we found was that:

- we paid a 3-5% penalty for the virtualization and resource management layers
- we increased the average core utilization from 15% to 30% (nearly a net 100% better utilization).

In Stage 2 we tuned each server to make each one 'cloud-ready'. To do this we needed to:

- strip out design assumptions that relate to hard dependencies relating to infrastructure (which architects/designers always deny exist, but are there — the worst of these were the designated storage point mounts)
- recompile the original code (because it was not usually multi-core ready or even Red Hat-safe) so that workloads could automatically be spread across multiple cores.

With all our organization's experience of doing this we could automate the second of these reasonably easily, using a multi-core aware compiler. The first requires people to go in and find those dependencies; there is no real way to automate this.

Today we have all the big scoring and batch applications running in our cloud. We have raised utilization to an average 60% (and sometimes as high as 80%), which represents an additional doubling of efficiency and is getting close to what we see on the mainframes. In effect we are now using 800 (1/3rd of the original number of) cores even after paying the virtualization and resource management 'tax'.

The pay backs after 2 years

The simplest way to describe the payback (because we do not discuss \$ numbers) is:

- we cut the number of physical servers by 300, from 1100 to 800
- we are running something like 3x the workload on those 25% fewer physical servers.

This will improve again. Eventually we will reduce the servers to 400 (from 800) for that same workload. But this can only happen when we possess complete confidence in the virtualization technology's ability to guarantee separation between the operational production workloads and the test environments: for the moment we keep these physically separated. In doing this we have already cut our capital compute budget by 50%. Our CEO's reaction was first — 'Wow' and then to ask 'why was it so high before?'

On the software side we have seen less significant savings — but remember that we were already a largely Linux environment (with some WinOS) and so had made many of the savings. In addition, we pay by the instance, and the number of instances running has gone up rather than down.

We have made further savings through standardization — what you would expect from simplification and removing oddities that cost us way more to support than the execution benefit. In addition, the total number of people needed to run the cloud has gone down significantly (although, in truth, we really did not have sufficient people before).

Lessons learned

For me this is pretty straightforward:

- first, make sure you have a credible business case before you start (and you had better hit the numbers you promise)
- second, it really does help to proceed incrementally when you go for a cloud route, to experiment up front before making major commitments
- third, you will only obtain the maximum benefit after you modify the application environment and code (so that these can take advantage of the cloud's capabilities); there is a big jump from just the initial virtualizing of an environment to what you

obtain after taking out those dependencies and recompiling.

In the latter case, I would not want to sound overly gloomy about this. We went from 15% to 30% utilization just by virtualizing (and most others would see at least an equivalent doubling of utilization — say from 10% to 20%). But the big kicker for us was going from 30% to between 60-80% after the modifications for the cloud were made. This is what delivered the big jump in value.

If you feel challenged by the first of my lessons learned, this was not as difficult as it might seem. Remember that we had much experience with the complexities of computing (we already had a sort of cloud — with a four-way Sysplex IBM mainframe, which we believe we can one day move towards our blade-based cloud). Suggesting the adoption of 'standardization, virtualization and automation' — if the commitment is carried through — builds a powerful business case. I was also confident, from previous consulting assignments, that we really could do better than before and that we had a good chance of meeting the improved price/performance objectives. Perhaps the hardest part was not sounding too confident initially, to the point that results were expected prematurely; all this takes time to execute.

Management conclusion

This financial information provider has implemented a private cloud that is producing ongoing and significantly improved price/performance in its production computing. It is achieving this through hardware savings as well as better use of people (the software savings had largely been made by adopting open source, like with Red Hat's Linux).

The keys to its success lie in the triple push for standardization, virtualization and automation. There was, however, the extra stage of improving applications to take advantage of the cloud multi-core environment — which doubled again the utilization.

Introducing a cloud may take some time, but it can clearly pay off in medium and large organizations. It can, however, only really pay off when there is true commitment to execution of the three principles. Too many organizations profess 'standardization' but fail to honor the concept. Failing at this stage likely will destroy a cloud's promise.

Value-based case study #5

Bluhalo: exploiting the cloud to cut costs and hugely increase capacity

Management introduction

Bluhalo is a full service digital marketing agency — based in Farnborough, England. It offers everything — from global digital strategy, social media, front-end website design to back-end website development, database design and development, multimedia, integration, security, usability and accessibility, along with web analytics. Its clients' websites receive over 5 billion page views every year and those clients include Liverpool Football Club, Tottenham Hotspur, the NFL in the UK, Blackberry/RIM, Andy Murray, MTV and others. Bluhalo is part of the GyroHSR Group and spearheads the Group's digital center of expertise.

Simon Green is in charge of all of IT for both internal (to support the business) and external (the infrastructure and web sites for clients) systems and support. He has been with Bluhalo for over 4 years. In this case study he discusses how, even for a relatively mature business (in digital terms), cloud computing is delivering significant operational and financial benefits.

The business problem

As part of Bluhalo's full service activities we are responsible for the design, build and then hosting and operation of web sites for clients. Our initiative to start to look at what cloud computing might mean for us was triggered quite simply by our concern about costs. Our Finance Director, for example, was not happy with the monthly outgoing costs against the derived income. The picture was not pretty; we knew we needed to do something.

At the same time I also was not happy from a technical point of view. There was too much that I felt we should be able to do fast — which we could not do because of the existing nature of our infrastructure. Before introducing our cloud we operated via a traditional managed hosting environment. As each project moved into production we would buy the server (or servers) needed not only to support day-to-day loads but also the anticipated peak loads. For some

of our clients — take Andy Murray being in the Australian Open Finals against Roger Federer — the peak loadings can be many, many times the average load. In effect we were putting in place peak-time resources dedicated to each 'site', which was clearly inefficient.

An illustration

While you can imagine why our Finance Director was less than happy, let me give you an illustration of what we have already achieved with our cloud approach. Then let me contrast this with what would have happened before.

In early February each year, the SuperBowl occurs. We host the NFL in the UK. During the SuperBowl this year, the NFL Digital Media Manager called me on the Sunday morning to say that his Forum was running slow (because of the rising level of interest) and to ask if there was anything we could do to improve that performance. Within minutes I was able to double the cores and memory dedicated to the SuperBowl Forum. The performance issue disappeared.

Before we introduced cloud computing, such actions might have taken 2-3 weeks to achieve the same effect. Why? Quite simply the way we, as a traditional managed environment working with a third party systems provider/manager, had to operate was operationally inflexible. Many signatures, authorizations and sign-offs were necessary before additional capacity could be brought online. In the NFL situation above, we simply could not have responded without weeks of notice — rather than a phone call on SuperBowl day and some straightforward actions.

Options considered

Before moving to our cloud, we had some 70 servers in-house for developers working on client sites and some 100 externally hosted servers on which our clients' web sites ran. In-house we used VMware's ESX for virtualizing our development server 'farm'.

This had proved itself time and again because ESX meant we could bring up or tear down servers in-house as and when we needed. I wanted this type of flexibility for our external operations.

We looked at many options for changing how Bluhalo worked. One option was, for example, to move to Amazon Web Services with EC2. We tried this with some pilot schemes — for example for video conversion — primarily to offload some heavy-lifting tasks to EC2. While these pilots worked quite well, we kept running into rigidities plus excessive SysAdmin overhead when Amazon Web Services hosts web sites. In addition there was a volatility associated with EC2 which did not suit our client needs. To make AWS and EC2 work we would have had to define, deploy and manage servers with a degree of involvement that did not fit our needs. What became also apparent about EC2 was that it was the application that mattered. In contrast with an ESX base and traditional servers, at least as far as applications are concerned, you do not have to build each application with EC2 in mind.

We looked at RightScale (at the suggestion of one of our tech-savvy clients). We investigated, Rackspace Mosso, ElasticHosts and others. A common failing was that we could not obtain the SLAs that are common with traditional hosting. We needed SLA security, because it is part of our service to our clients that they know what they are obtaining from Bluhalo when we are running their sites.

Then, almost by accident (a sales rep. made a cold call to me at the right moment) we came across DediPower. In brief DediPower is a hosting company that runs datacenters. It was reasonably local (in the UK) plus it offered an international reach (UK, USA, Hong Kong, etc). It was also about the same size and age as Bluhalo — which mattered to me: I explicitly had always said I did not want a provider that would tie us up in red tape and bureaucracy, as so many of the larger hosting companies do — we need to get things done. Yet any provider also had to be big enough to support and guarantee credible SLAs. Finally, DediPower turned out to be willing. One reason we chose DediPower was that its people recognized from our first meeting what our needs meant — and were willing to work with us to satisfy these.

Dreaming of a dream environment

At the start DediPower was not offering us a cloud (in

fact it was expecting to offer us straightforward managed hosting). While it had been involved in a small number of virtualization projects, it had not taken the next step or moved up to the large scale.

We described our ‘dream environment’ — one where:

- the overall cost would be less than our existing solutions yet have at least 30% more capacity
- everything would be billed paid for by the hour
- provisioning would take minutes or seconds — not weeks (and without copious paperwork)
- we could easily move web sites from development into production
- interactions between Bluhalo and our chosen provider could be on a ‘both sides know each other’ basis
- we could leave the buying, installation and implementation of the hardware, communications and operations to those who specialize in it (which we did not want to have to do).

DediPower went away, thought and came back to us with a VMware approach (it had experience of Xen as well) that on close examination was greatly appealing. Once we had agreed what was involved, it took responsibility for the acquiring and providing the hardware, communications and operations infrastructure. (Since we went live, DediPower now offers more or less this same cloud/virtualization package to others.)

The DediPower proposal was based around a resilient, rack-based processor solution with a SAN attached, designed and implemented by DediPower. We originally had about 100 physical production servers. We now have the ability to run all those 100 servers (though not all of those 100 always ran all the time) on 9 physical machines operating as an ESX cluster with the SAN attached.

Measured solely by the reduction in power consumption, the gains are astonishing. In addition we moved everything to modern hardware (some of our existing servers had still been running on elderly Pentium 3 servers).

We also obtained much more out of this architecture than we had had before. We were looking for 30%

more capacity with improved resiliency. We are actually utilizing only 10-15% of the platform at normal levels; thus we have over twice the excess for client expansion and peak loads than we had sought — all within a price that was around 25% lower than what we had been paying.

In addition, because of the nature of the contract with DediPower we have full access to vCenter. This means that we can bring up and take down servers as our client needs require. All we need do is liaise with DediPower to enable the appropriate routing.

In effect, DediPower now looks after our production infrastructure for us. As long as its people are kept informed, the DediPower attitude is that it should be like just an extension of our IT department which happens to handle all the monitoring, backup and other operational characteristics necessary. This 24x7 support is a blessing for us. We knew we could not afford the quality of SysAdmin staff that DediPower can provide. After all, we are a Web company; we are not an infrastructure company. That is DediPower's speciality.

Making the switchover

Extraordinarily for the IT world, the switchover from our old hosting to the new was smooth. As I recall we only had a couple of sites which had any noticeable downtime and this was only due to external influences. We established VPN tunnels between the old hosts and the new and migrated across.

Even better, we were able to use the switchover as an opportunity to accomplish a substantial degree of rationalization. Looking back, that makes our smooth switchover perhaps a little more incredible. However we did take our time. We did not try to do everything at once. We moved steadily over 2 months, client by client. We only went live with each site in the new cloud once each rationalized environment was ready and tested.

In that housecleaning we did not virtualize what had existed. Instead we used the opportunity before each transfer to update to the most recent versions of PHP, Apache, other utilities and even operating system (we use Linux and PHP almost exclusively on our servers). This is the sort of work that tends not to happen. We used our transfer window to perform a thorough job so that what went into our new 'cloud' was 'cloud-appropriate' for the future. In effect we created brand

new OS virtual machines, installed the latest versions of software plus our applications, tested these and copied them to the cloud.

An important aspect of this was, and is, facilitated, by ESX. To move from development into production we just move a copy of the developed (and tested) virtualized server from our internal environment to run identically in the ESX cloud in the DediPower facility. Apart from the time to move the image, it is as simple as this and, as we have big pipes between our offices and the DediPower data centers, it does not take long.

Best practices and lessons learned

My number one best practice, the fruit of lessons learned multiple times, is that having good templates for servers for VMware makes life much easier and infinitely more predictable. We make sure we have rock solid versions of the OS with other necessary software all working for the various scenarios that we think we will encounter with our clients. We create templates for each of these scenarios as individual VMs. When we need to bring up a new server of a particular kind, we use the appropriate template and go. This saves large amounts of valuable time, both in implementation and in avoiding building and 'discussing' what should be incorporated.

We treat virtual servers as if they are real servers, not just candied abstractions for endless consumption. If you permit the latter, a mess will result. It is too easy to fall into the mindset of a 'virtual server here, a virtual server there' (because they seem to cost almost nothing) and end up with undocumented and all too often unrecoverable chaos. By treating every virtual server as real, we maintain discipline.

My third best practice goes alongside this. One of my intentions has always been to avoid VMware 'spread', which is never pretty. To avoid this we present ESX internally as if it is a service running servers; development people request an image (copied from the templates) for particular authorized tasks. In this way our application developers are not seduced into flexing ESX (for example requesting or building servers which then may or may not be used and backed up, etc). We categorically separate the function of the developers (who know how to build beautiful web sites with all their relevant tools) from our cloud infrastructure specialists (who know ESX and how to manage this). This works for us because the developers do not

need to understand how ESX functions, only how to exploit what can run on it.

One other best practice we have found works is to have a good IP allocation map. Though this is little more than good planning, it does make implementation and operation that much more straightforward.

Moving forward

Thus far I have focused mostly on our external hosting, with occasional references to relevant internal practices. Would we think of moving our internal systems into a cloud environment such as we have with running DediPower? Absolutely, yes — based on our experience so far.

What is holding us back? Not the communications — we have 100MB links between our offices and to DediPower as well as using DediPower's New York and Reading data centers. In addition, we have already spotted significant amounts of elderly equipment that we believe could be retired from our in-house infrastructure whose function could happily and reliably exist in our cloud: this would further cut day-to-day running costs as well as minimize capital expenditures. We envisage running an Exchange Server here clustered with an Exchange Server in the US running on the same VMware backbone so that wherever you are working, you will be supported from the nearest/fastest location.

Next on the to-do list is to move everything (except, perhaps the file servers — simply to avoid connection bottlenecks) across into an equivalent, internal

Bluhalo cloud. We have not had any connectivity problems and no downtime since we moved into our external cloud. We have less management involvement than we had before, and with better as well as crisper results.

Finally we are planning to merge the management of our internal and external VMWare deployments into a single vCenter setup, allowing centralized management of both and live migration of servers between physical locations (eg Farnborough office to London office).

For me, moving into the cloud made — and still makes — complete sense. I am just looking forward to making the transition with our remaining equipment, and saving more costs.

Management conclusion

There are many who believe that cloud computing is either only for start-up like organizations (where no IT infrastructure exists and the cloud takes its place) or for large organizations who have the ability to manage whatever they require from a cloud. Bluhalo is a proven example of how an established, though relatively small, business can exploit a cloud.

What should be apparent is how much can be achieved. Bluhalo obtained a 30% reduced cost base supporting the same number of sites but with 80-90% additional capacity and flexibility, and increased resiliency. That should impress any size of business. Internal cloud possibilities are enormous, for those prepared to take the initiative.

Value-based case study #6

D&B innovates by integrating coAction's external SaaS platform with its internal applications to create a new customer Collections product

Management introduction

The risk management side of Dun & Bradstreet (D&B) has traditionally helped commercial customers to make informed credit decisions. It achieves this by collecting and then providing information about the creditworthiness of buyer organizations in order that goods and services — not only to prevent customers from providing goods/services to organizations with low creditworthiness, but also to enable customers to differentiate the levels of creditworthiness within their portfolio, and offer appropriate credit terms to each customer.

The approach has worked for many decades, with constant improvement. But, while the D&B Interactive (DNBi) service — one of those improvements — has extended the ability to provide online decision assistance about whether or not to grant credit, as well as analytical tools to continuously monitor and adjust credit terms, it (DNBi) did not offer support for making sure that payments for Receivables were collected once invoices had gone out. DNBi decided to add a software service to improve its customers' abilities to collect payments.

In this case study, Jess Urriola — who is the Manager responsible — discusses what DNBi implemented and why. What is intriguing is how DNBi has melded its in-house credit platform with an external SaaS platform (produced by coAction) to deliver in months a receivables management service for DNBi's customers. SaaS tends to be thought of as 'being wholly out-of-house'; but, as DNBi has shown, it is quite feasible to create a combination of in-house and out-of-house. In so doing DNBi has implemented something original, investment-saving, scalable and do-able in an extremely short period.

Starting out

Our core business at DNBi has been the delivery of information to our customers about the credit worthiness of third parties, particularly when making sales. Using DNBi's online (web-based) services along with

the data backing those services, customers are able to take decisions about whether to accept an order or not, and on what credit terms. However, once the decision to grant credit had been made, we did not have an offering which closed the loop in the 'Order-to-Cash' cycle, one that took the process through from the raising of an invoice to receiving payment for it (so-called Receivables).

When we looked at the market place we found that, at the top end, the major enterprise 'suites' (SAP, Oracle/PeopleSoft/etc.) integrated the receivables function into the myriad of other application modules that make up ERP. But, if you were not a user of such enterprise applications, you had a myriad of possible receivables providers — none of whom enjoyed significant market share (the largest had revenues of about US\$10M/year).

With D&B's 100,000s of customers, we saw a significant opportunity to extend our business. We decided that this 'receivables management' was an area where our existing information together with added function could be leveraged for our customers' benefit which we would call DNBi Collections Manager. Our target market initially was what we call the commercial mid-market: 200-1000 employee-type companies in the trade credit space. (Subsequently we have produced a scaled-down edition of what is now sold as DNBi Collections Manager –Professional Edition, for smaller businesses in September 2009, and we are working on a scaled up Enterprise version, for Fortune 500 size companies for launch later in 2010.)

Creating the Collection Manager product

In late 2007 and early 2008 we sat down to examine what it would take to build an in-house solution, much as we had done for most of our existing DNBi services. Pretty quickly we decided that we could afford neither the cost nor the time to build in-house (we wanted to enter the market in 2009). We concluded it would be faster and better either to buy an

existing receivables management vendor or to co-operate with a partner — assuming that we could find a 3rd-party application that would integrate with our existing DNBI creditworthiness platform.

We started initially by going down the acquisition route in 2008. For various reasons this did not work out as we had hoped. With that no longer viable we decided to look at some form of co-operation or partnership with a third party vendor.

Remember, and this was another key element in our approach, we approached our decision processes always looking for an already working receivables management platform which had the flexibility to combine with our existing DNBI offerings. We were not looking for a technological platform on which to build a receivables management solution — but a receivables management solution that already sat on a well thought-through platform.

Choosing our partner

As I alluded to above, the receivables management 'market' is diverse with no dominant players. After due diligence, applied to both the receivables management functionality and the underlying technology platforms, we decided to partner with coAction.

We decided this for a number of reasons. coAction was local to us in New Jersey which made for much simpler working together. But three other reasons mattered more:

- from a receivables feature/function standpoint, we felt the coAction offering addressed the capabilities that we wanted to offer to DNBI's customers — coAction specifically addressed the collection team's needs and its management (rather than just an individual collection agent's needs)
- coAction offered a platform that was web-delivered; this was consistent with where we wanted to go technologically, particularly given its support for ease of introduction, adaptation and then maintenance from our (DNBI's) customers' perspectives
- scalability came as a given, because coAction is offered on a Software as a Service (SaaS) basis; we could start small and grow without having to incur significant capital or asset investments before we had our customers.

Furthermore, this approach had the significant benefit that it was faster and less expensive for us to create an offering for our customers. Plus, doing it the SaaS way enabled us to enter the market with a much less expensive initial service than would have been possible using the more traditional customer-hosting model.

Just to offer one insight into what this meant: in early 2009 we signed the agreement with coAction. In April 2009 we launched our new product as Collection Manager for DNBI which is now sold as an add-on module to the DNBI web-based interactive risk management platform. We could not have done this so fast if we had tried to build it in-house, nor as cost effectively

The architecture

We took the coAction platform and its existing receivables/collections workflow solution. This was already an excellent tool in that it enabled customers to aggregate all key data — whether customer ageing data, customer contact data, status history, notes about and records of conversations with customers, etc. — with excellent communications tools. For example, one can email from within the coAction platform (no need to exit out to an email client); there are also call-scripts already incorporated.

Perhaps more importantly, there is an embedded automation engine. It is this that enables customers to create rules which segment customers who owe money into different bands. Each band can then have unique workflows suitable for that particular type of situation. As we had predicted, there are significant efficiency gains in doing this — fewer people can do more as well as be more targeted.

Where we believe we go a stage further in the architecture is that we integrate D&B scoring into the Collections Manager. We use credit scores in the Collections Manager strategy engine to work out priorities so that receivables personnel focus their attention on those customers which our data indicates as showing signs of being potential defaulters. By exploiting the vast amount of existing DNBI data, our Collection Manager customers can then concentrate their collection efforts on where it matters.

Other collection tools may have strategy engines. But mostly these use 'ageing' or 'past due days' or similar data to establish the collection strategy. Where DNBI

differs is that we bring a broad view of each customer's risk — through exploiting our real time credit bureau data. This is what we mean when we talk of risk-based collections prioritization.

An example of risk-based collections prioritization

Let me offer an example. A DNBI customer may have (say) 20 invoices of around US\$25K each outstanding (say from 20 different companies). Traditionally, collections management will determine which invoices are oldest and most overdue — and start focusing on collecting the outstanding payment from these. Or the A-Z or Z-A approach may be taken — when collections agents start at A and work to Z, or vice versa.

But what if 3 of those 20 are showing one of the patterns that D&B's analysis says is an indicator of future payment problems (D&B possesses analyses that can do this)? For our customers, the collection priority needs to be focused on the invoices from the 3 customers where D&B's real time data suggests problems may occur. Doing this is in direct contrast to the approach warranted for (say) 4 of the other customers — all of which have rock solid payment records (with other suppliers as well, as D&B knows) even if they do have a history of being a little behind in making payments; a reminder email may all that be necessary. Understanding this and acting appropriately can be the difference between being paid or receiving nothing.

This focuses the collection of receivables on where the payment risk lies. It makes best use of scarce staff in the receivables department. It brings in the cash that matters.

Our experience to date

In practice the way we are working with coAction is that the Collection Manager runs on its platform on its boxes. We (DNBI):

- buy this as a SaaS service which we have integrated with our web-based online DNBI platform
- sell the resulting combined functionality as a service to DNBI customers who see a seamless service over the Internet, even though the Collection Manager is being hosted and delivered from the coAction platform.

It turns out that the receivables function in most organizations is quite different to the credit analysis function. We had been used for years to selling upgrades or enhanced capabilities to the credit function people; they understood what credit scoring and credit data meant and were familiar with D&B as a provider. This was not true of the receivables people. To them the concept of having credit scoring and history data about the organizations they had to collect payments from was novel. In effect we had to make a whole new sale to the decision makers responsible for receivables — who were not necessarily familiar with what D&B could provide to their function.

There was also a learning curve aspect — for both our DNBI users as well as potential customers. Many receivables functions did not see that introducing some degree of automation — via a receivables application — could improve payments and reduce outstanding debtors. Similarly, the idea of using our credit score data as a predictive tool was novel. There were, therefore, initial obstacles — from both sides.

That said, once they understood this aspect plus they accepted the notion that Collections Management was a responsibility that would benefit from the application of automation combined with information, then we started motoring.

Once Collections Manager is installed the feedback has been good. Customers like the capabilities to send emails and to call customers from within the coAction platform; even better, in doing so, the tool possesses an 'Actions Log' that tracks every activity associated with a customer account down to the invoice level — and this is date and time-stamped. Thus, if a customer says on a subsequent call 'oh you did not call or email me' the collection agent can immediately see when a conversation was held or an email was sent along with what was discussed (even if this had been conducted by another agent).

From a manager's perspective, besides the automation engine prioritizing where attention should be focused, it is possible to see which collection agents have done what. Plus there is evidence for compliance and auditing.

What is different is that everything is brought together. Before a collection agent might be using the Accounts Receivable application to find the data about outstanding invoices, then output this to a spreadsheet, use an email client for email, open a

customer contact application to find whom to call as well as to make notes, and use the phone for calls, etc. — with minimal integration (and with no access to external real time risk data). When you are bouncing in and out of applications like this, it is difficult to keep track of, or organize, all the data — and even harder to manage effectively the receivables function.

With Collection Manager on the coAction platform, integration is grouped together. This is then delivered through a customized interface, available in the coAction platform to comply with D&B branding. All is brought together in one 'solution'. Everything changes.

Lessons learned and best practices

One of the areas where we struggled, and where I suspect others struggle, was in over-projecting existing brand strengths into new product areas. In our case we thought that the D&B name in credit worthiness and scoring would possess an equal impact with the receivables community. In this we were overly optimistic. We found we had to educate receivables people about how automating the support receivables and integrating real time credit information can change the priorities in which receivables management approaches the selection of which payments you need to pursue (as my example described earlier).

The original Collection Manager platform was designed to manage open invoices. But this is not the whole collections story. There are credit notes and adjustments/amendments and partial payments. The data for this exists in the Accounts Receivables application and we have had to work hard to extend the Collection Manager for some customers to accommodate this. The lesson learned here is that you cannot always just simplify, even if in that simplification there is real value.

Related to this is that customers do not always have to hand the 'cleansed' data — even for invoices — that they imagine. As often as not we have had to coach customers on what 'good data' looks like for exploitation of the Collection Manager platform in

order that collections can work effectively. Because of this the first few customers took longer than we expected to be up and running — 90-120 days. We have learned our way through this and we are heading to meet our original goals of 30-60 days, because we now know what needs to be done.

Within the Collection Manager's automation engine we incorporate collection strategies; these provide pre-canned D&B recommended, and proven, approaches for how collections managers should think about, segment and handle high risk customers, increased risk ones, low risk ones, etc. One best practice that has emerged is to use a customer's own D&B-based custom scorecard for the priority weighting to be applied to receivables invoices. If, however, a customer does not have such a scorecard, then the next best is to take decisions using the D&B Commercial Credit Score; this is a predictive indicator of the likelihood of a customer paying late over the next 90 days, one which has been proven to work.

I should add one other area where we have confirmed a lesson — working with a responsive partner makes a big difference. coAction has been excellent in the way it has co-operated to support us and our requirements. Our decision to go with a SaaS platform has been wholly vindicated. When we have needed changes or additions, it has worked hard, almost always with mutually beneficial results. Indeed we have seen improvements in the underlying coAction platform arrive throughout the past months. As the business scales up, we have confidence that this responsiveness will continue.

Management conclusion

For those who might have thought that choosing an SaaS solution is black or white (as in either externalized or not), DNB's Collections Manager with its integration of D&B internal resources and coAction's rule-based platform is incontrovertible proof of what is possible.

That it could be up and running and offered to customers so fast is another striking aspect.

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

All rights reserved; INSIGHT-SPECTRA may be reproduced for individual use. However, bulk reproduction (more than 5 reproductions -- whether by electronic means or in print) requires prior written permission from the Publisher, C3B Consulting Ltd.

© 2010 C3B Consulting Ltd.