Business Assurance Requirements for Customers in a Services Marketplace

Bryan Stephenson, Jun Li, Sharad Singhal

HP Laboratories
HPL-2008-169

Keyword(s):
cloud computing, business assurance

Abstract:
The Mercado project explores a services marketplace that provides a business-over-the-Internet environment, where a business user can specify business tasks and run a business lifecycle using services provided over the Internet. Our objective is to build business assurance techniques within Mercado so that customers can be confident that it is safe to make the success of their business dependent on such marketplaces and services. To uncover requirements on data exchanged in such marketplaces, we perform a threat analysis on a fictional business that runs a marketing campaign using online business services. We detail specific customer requirements and research areas uncovered by the identified risks. We position these research areas in a service environment that significantly differs from traditional enterprise IT environments.

External Posting Date: October 21, 2008 [Fulltext] Approved for External Publication
Internal Posting Date: October 21, 2008 [Fulltext]

Business Assurance Requirements for Customers in a Services Marketplace

Bryan Stephenson, Jun Li and Sharad Singhal
Hewlett-Packard Laboratories
Palo Alto, California

Abstract

The Mercado project explores a *services marketplace* that provides a business-over-the-Internet environment, where a business user can specify business tasks and run a business lifecycle using services provided over the Internet. Our objective is to build business assurance techniques within Mercado so that customers can be confident that it is safe to make the success of their business dependent on such marketplaces and services. To uncover requirements on data exchanged in such marketplaces, we perform a threat analysis on a fictional business that runs a marketing campaign using online business services. We detail specific customer requirements and research areas uncovered by the identified risks. We position these research areas in a service environment that significantly differs from traditional enterprise IT environments.

1 Introduction

The Mercado project explores a *services marketplace* that provide a business-over-the-Internet environment, where a business user can specify business tasks and run a business lifecycle using services provided from service providers over the Internet [1]. Our objective is to build business assurance techniques for Mercado so that customers can be confident that it is safe to make the success of their business dependent on services delivered over the Internet. Service providers can register and advertise their service capabilities at the Mercado portal, and service consumers can discover, select, and consume these services based upon their unique business requirements with Mercado acting as the broker. A business application is composed of many such selected external services. We hypothesize a services marketplace accessed via a Mercado portal running what we call a *Business Operating Environment* (BOE). The core functionality provided by the BOE includes:

- Requirements engineering and business process definition: Requirements are captured, refined and mapped to candidate services.
- Service selection assisted with aggregated web information: The particular services and service providers are selected. People are likely involved to guide and validate selections and make final choices.
- Service-oriented architecture: Runtime support provides the capabilities to discover, aggregate, and invoke services, and monitor service execution against high-level business objectives.
- Cross-organizational collaboration: Service providers, customers and business consultants can coordinate actions, share knowledge, and establish electronic service contracts.

Our main objective is to uncover and understand from the perspective of the customer some of the new issues and requirements likely to emerge for a business using business-level services from such a BOE. To identify such issues, in Section 2 we present a use case using a fictitious business called Nullco that runs a product marketing campaign using online business services and a hypothetical Mercado portal. In Section 3 we present selected results of a threat analysis of this use case using the Octave-S methodology [2] as a means to extract these requirements. Section 4 details how the identified risks and countermeasures imply specific customer requirements for a BOE to provide business assurance. Section

5 discusses the research areas uncovered, and Section 6 identifies the work that other people have undertaken in the related research areas. Section 7 presents some conclusions from our work.

2 Example Use Case Description

The particular use case we investigated was the Vice President of Sales for Nullco initiating, and Nullco's marketing people executing, a new marketing campaign for a new product, using many online services that can be discovered and then selected via the Mercado portal. We made several assumptions for this use case:

1. Nullco is a small business of around 100 employees, $10M yearly revenue, 10% profit margin.
2. Nullco can export marketing campaign information to contracted service providers using web services capabilities.
3. Contracted service providers can export campaign results (like qualified responses to the campaign) to Nullco's in-house CRM system using web services capabilities.
4. Nullco has a few internal IT systems, most notably their in-house CRM system, but Nullco mostly uses external service providers to meet its business needs which have a significant IT component.

Nullco's Vice President of Sales and Marketing meets with the marketing department staff to plan and launch the marketing campaign at a high level. The schedule, goals and budget are determined, as shown in Table 1. After this initial meeting, the marketing department runs the campaign with periodic reports to the VP on progress.

Table 1: Criteria Defined for Product Marketing Campaign Use Case

Goal in context:	The successful completion of a product marketing campaign is measured by the number of qualified leads generated per unit cost
Scope:	This use case starts when the VP of Sales and Marketing decides to run a campaign until the time the final report is produced showing the number of qualified leads generated.
Pre-Conditions and Assumptions:	A product exists. Product information and descriptions exist. Existing customer data is available to determine demographics of those who already purchased the product. A broad number of services offered on the internet are used to execute the use case. Marketing constraints: budget = $50,000, Schedule = 3 months to meet marketing event; and target to result in a 10% increase in sales.
Success End Condition:	The campaign material is created. The campaign is successfully completed including delivery of offers, and a final report delivered to the sponsor. Goals met or exceeded: (a) the number and value of qualified leads, (b) % of leads accepted by the sales organization, (c) incremental revenue per dollar of marketing campaign cost (ROI)
Failed End Conditions:	1. A step along the way fails and no recovery path is available to complete the campaign. 2. A low number of leads results from the campaign.

The campaign is run by using many services provided by many different service providers, as shown in the overview of campaign execution in Figure 1. A data mining service (provided by SuperAcim) is used to identify the targets for the campaign. A creative agency is used to design various campaign materials, including banner ads with text, graphics, and video; brochures to mail; email messages; and

a web landing page for potential new leads to register with Nullco. The campaign materials are hosted in the content management service provided by SuperWun. Next, the campaign is launched with targeted printed brouchures through direct mailing, email, and banner ads. A campaign tracking service is used to gage effectiveness and fine-tune the running campaign. Finally at the end of the campaign, the generated leads from the multiple campaign channels are stored in the CRM system of Nullco.

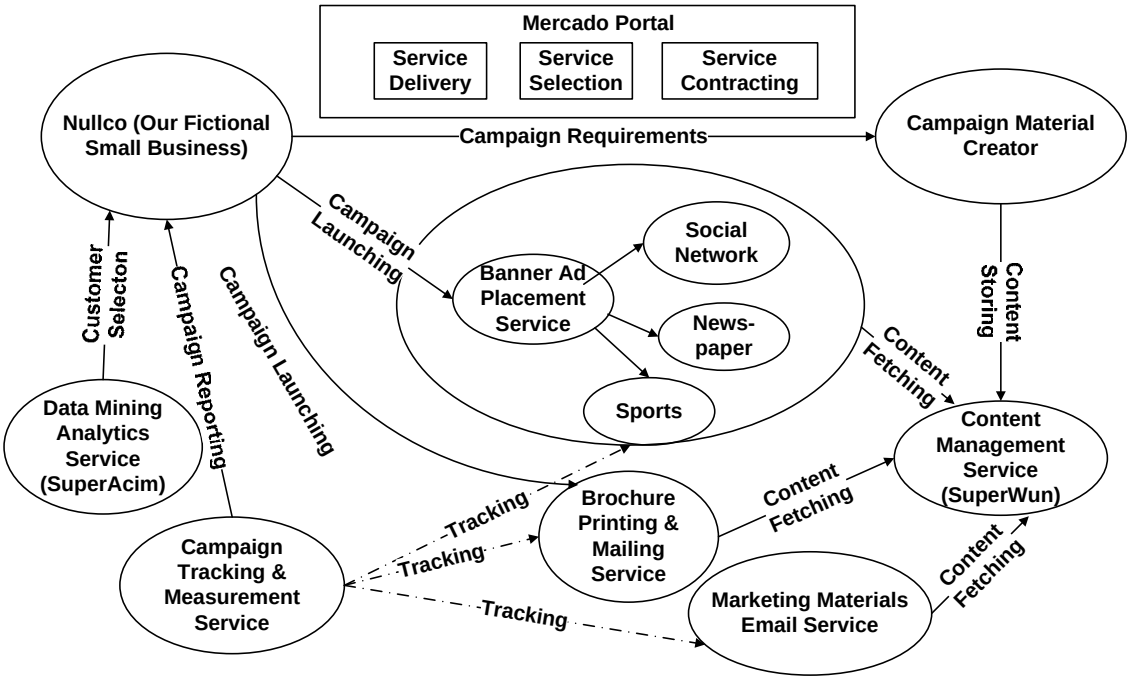

Figure 1: Services Involved in Product Marketing Campaign Use Case Example

3 Threat Analysis of Example Use Case

We use the Octave-S [2] methodology to conduct a threat analysis of the example described in Section 2. This methodology was selected because it is tailored to enable a small business which outsources most or all IT functions to perform a self-assessment. We focus on only a subset of the analysis pieces recommended by the methodology which were useful for this purpose, which includes identifying assets, prioritizing the five most critical assets, and pinpointing threats to those assets from the various sources. Based on the above constraints, the five most critical assets identified in this exercise were:

1. Nullco's contracts (both physical and digital) with service providers.
2. Nullco's internal CRM system holding campaign results information: responses, leads, opportunities, campaign-generated revenue, post-campaign analysis, etc.
3. External content repository at SuperWun that holds Nullco's creative assets: campaign brief, email templates, postal mailer designs, telemarketing scripts, banner ad designs, white paper offers, etc.
4. The customer data provided to SuperAcim for it to provide candidate leads for the product campaign.
5. The Mercado portal which Nullco uses to locate and conduct business with service providers.

The raw materials from the threat analysis, including the complete list of threats, are available from us in [3]. We summarize below the most interesting threats for our purposes. Each threat is accompanied with a corresponding countermeasure.

3.1 Selected Threats to asset: Nullco's contracts with service providers

Threats from accidental causes
Risk 1: Several versions of electronic or paper contracts are exchanged between Nullco and another company. Nullco verbally agrees on version N of the contract, but due to a mistake, both companies sign version N-1 of the contract. The other company is unscrupulous and holds Nullco to the signed version N-1 of the contract, to its detriment. Countermeasure: Process control with approval process on both sides.
Risk 2:. Someone mistakenly puts a CD with a copy of a contract into their bag to take home. Their child grabs it and it ends up being inappropriately disclosed, for example posted to MySpace. Due to the many internet archives and caches, it is very difficult and expensive to remove this data from the internet. Countermeasure: Encryption of sensitive data
Threats from deliberate causes
Risk 3: Nullco has a company policy which requires both legal review and approval, and the company president's approval before signing an electronic contract above $50,000. However, this is a paper-based policy control only; no technical control is in place to enforce the policy. Because this process is not enforced in a real-time manner, someone who is authorized to electronically signed contracts (and has the credentials to do so) signs a $500,000 contract without getting the required approvals. The contract is not beneficial for Nullco, but is not so heinous as to be obviously fraudulent (a jury could be convinced that Nullco did intend to sign it and then reconsidered). Since electronic signatures are legally valid in the US, Nullco may be obligated to honor this contract despite its employee violating company policy. Nullco might be able to convince a jury that the signature should be invalidated since it violated its company policy, but case law for the jurisdiction in the USA is untested in this scenario, and Nullco may not want to be the first to test it. Countermeasure: Process control with approval process, and fine-grained access control that takes the approval results into account for contract signing.
Risk 4: An officer of Nullco who is authorized to make deals and sign contracts is close to completing negotiations on a contract with another company, and they decide to sweeten the deal so the other company will sign it today, in time for this result to be considered for the officer's imminent performance review. If they would have waited it out and ignored the personal incentive to push for faster signing, Nullco would be better off because it wouldn't have needed to sweeten the deal. Countermeasure: Process control with approvals and continuous monitoring of business activity for abnormal patterns.

3.2 Selected Threats to asset: Internal CRM system holding campaign results

Threats from accidental causes
Risk 5: The data analysis company SuperAcim accidentally releases to Nullco its internal information too broadly, which includes knowledge from Nullco's competitors. Nullco imports the knowledge into its CRM system, including information about a competitor that it shouldn't have. Nobody in Nullco detects this situation. Later, for an unrelated reason, its CRM database is subject to legal discovery. When this happens, it is revealed that Nullco holds information about a competitor (such as some of their private financial numbers, forecasts, or customer contacts) which Nullco shouldn't have and didn't know about. This causes damage to Nullco's reputation, even if Nullco is not at fault. Countermeasure: Data quality validation. This is an example of a class of problem which will happen more frequently as larger amounts of data are sent between different companies with lower amounts of human oversight and the IT supply chain disaggregates and gets very long. A mistake by one person in one company can jeopardize the business and reputation of another company *despite all involved trying to do the right thing*.
Risk 6: Some trusted business partners have access to Nullco's CRM system in order to streamline processes and reduce costs. Personnel at these partners sometimes keep local copies of some of the information from the CRM database in order to provide their services, especially when they are traveling. It is possible that such information stored on a smartphone's MicroSD card could be misplaced at a conference room, stolen, and then sold to competitors. This could result in the competitors stealing Nullco's leads or customers, or Nullco needing to notify customers that their personal information was breached. Countermeasure: A well-defined data management policy, with a data synchronization package installed at the client devices to perform necessary data encryption.
Threats from deliberate causes
Risk 7: Nullco gets web site tracking data per click, when cookies are enabled per user, from ClickMonitor each day. At the moment, It is 1 GB of data per day. Nullco has no way to verify if this data is real or not because no model exists. So ClickMonitor could be randomizing previous data and changing timestamps so they can do much less work and still get paid. Nullco may be making critical business decisions on incorrect data. Countermeasures: Continuous monitoring and data quality validation
Risk 8: A telemarketer may be including false responses to Nullco. For example, the telemarketer didn't call or the contacted person wasn't interested and didn't want to be contacted, but the telemarketer said they did want to be contacted so they could claim another response to the campaign. This hurts Nullco's image when Nullco calls them later and wastes money on false leads. Countermeasure: Process control that adopts data quality validation, e.g., to compare the written script of the telemarketer to some call recordings of the phone conversations.
Risk 9: An outsider with authorized access to Nullco's CRM system alters or destroys some of Nullco's data. After this happens, they deny that they did it. The motive could be they are late delivering something they agreed to deliver by a certain date, and they sabotage the data in the system to try to buy themselves some more time by using this as an excuse for being late. They try to plausibly deny that they were the one who sabotaged the system by claiming that a hacker must have gained access to their PC and impersonated them to access the system and from there deleted the data. In reality, they sabotaged the CRM system and ran anti-forensics tools on their PC to cover their tracks and gain the ability to plausibly deny being at fault. Countermeasures: business transparency and continuous monitoring

Threats from System Problems
Risk 10: Some important customer data lives in a legacy CRM system on a very old machine. This machine crashes with an unrecoverable hardware failure. Nullco has a backup tape of the CRM database, but this is a proprietary format and can only be read by a very old version of the CRM software. This software can only run on a very old version of the OS for which support has been discontinued. Furthermore, this OS version can only run on very old machines, which haven't been for sale for several years. These aren't even for sale on Ebay anymore, so Nullco has no way to recover this data. Countermeasure: A backup service equipped with correct data interpreter, not just simple raw data storage.
Risk 11: In preparation for the end of the quarter push, several sales people perform a bulk transfer of information from the CRM database to their laptop to process it over the weekend. A stealthy "end of the quarter bit twiddling" virus on 90% of their laptops modifies several records in this data without being detected. When they get back in the office the automatic data synchronization tool imports these modified records into Nullco's database. By the time this is detected things are pretty messed up and Nullco needs to load a backup copy of the CRM database, causing Nullco to lose 5 days of work at a bad time and miss its finannical numbers this quarter. Countermeasure: no clear solutions devised

3.3 Selected Threats to asset: External content repository at SuperWun Corporation

Threats from Natural Disasters
Risk 12: The telemarketer relies on real-time access to Nullco's content repository to get their job done. Due to a natural disaster, SuperWun cannot provide the real-time access when required. The contract with SuperWun has an excellent SLA acceptable to Nullco, including a 99.999% availability guarantee or else the service is free and they pay a substantial penalty. However, the contract relieves them from paying any penalty or consequential damages resulting from natural disasters. Due to their high availability guarantee with substantial penalties, Nullco felt safe entering into a contract with a very inexpensive telemarketer which requires paying for their staff hours ($50,000/day) when the service they depend on is unavailable for any reason. Thus, Nullco is liable to pay the telemarketer even though they aren't getting any work done. This is an example of a class of problem which results from the dependencies between services and the variety of SLAs agreed with the services, and the inability of a service consumer to accurately estimate the risk of downtime of a service provider. A natural disaster in a remote location can impact the business, even if the business does not operate there. Countermeasure: A rigorous data center availability certification process.
Threats from accidental causes
Risk 13: Someone drags the wrong file to the right place, or the right file to the wrong place, and corrupts some data in the external repository, such as the banner ad designs. The next day, this is discovered and the correct banner ad designs are uploaded, but the banner ad service has already retrieved the wrong designs and sent them to 30 subcontractors in 30 different countries. They charge Nullco $1000 per subcontractor to manually intervene and replace the ad designs. The wrong ad design gets run for several days in some of these countries. Countermeasure: Process control that adopts data quality management

Risk 14: A design firm creates 20 different banner ad designs for Nullco. They notify Nullco via email that the designs are ready for review on their web site and will be pushed into the SuperWun repository in two days time, but Nullco doesn't receive this notification. Since the design firm doesn't hear from Nullco, they push the banner ad designs into the repository, and the banner ad service retrieves them and serves them per Nullco's contract. However, two of the banner ads don't meet Nulco's brand standards and one is offensive in Germany but is run there anyway, damaging Nullco's reputation and brand image. Nullco needs a stronger approval process to review the banner ad designs before they are served. Countermeasure: Process control that adopts approval process
Risk 15: Nullco contracts with SuperWun to host up to 10 GB of data per campaign, which is normally plenty of space. An agency uploads some video content which is very large and fills the quota. This prevents other agencies from uploading their content, for example preventing updated banner ads from being pushed into the repository, resulting in using the old banner ads instead of the new ones which makes the campaign less effective. Countermeasure: Continous monitoring and fine-grained access control that takes into account the file system quota as the attribute for access decision
Threats from deliberate causes
Risk 16: Copyrighted material (such as a photograph) is maliciously or accidentally attached to some of Nullco's banner ad designs. This is served to 5000 internet users by Nullco's contracted banner ad service. The copyright owner sues Nullco and the banner ad company for $100,000 per banner ad served, or $500M. No one can figure out who attached the copyrighted material. A jury could find that Nullco is at fault and fine Nullco enough to bankrupt it. Countermeasure: Process control that adopts data quality validation, along with with approval process Notes: Technology exists to sign data to prevent insertion of unauthorized data, or detect who added something like a copyrighted image to a dataset. The issue is that this technology is hard to use. There is an opportunity to integrate these capabilities into a solution such that they become seamless, or at least easy to use.
Risk 17: A Nullco competitor uses social engineering to trick a SuperWun employee into divulging protected company information. Example: they pretend to be the contracted printing service company for the brochures to be printed for the campaign, and falsely claim that their web proxy is down but they can send and receive email. They send email and convince the SuperWun employee that they need the brochure design files by 2:00 today in order to make the deadline for printing the brochures. The SuperWun employee is suspicious of this, so the attacker offers to also telephone them so that SuperWun can verify their caller ID. The SuperWun employee doesn't realize how easy it is to fake caller ID, so after they receive the telephone call with the printing service company's business name and phone number, they verify that the phone number is correct in their phone book. Then they send the brochure in email to Nullco's competitor, who sees it weeks before they otherwise would and gains unfair advantage. Countermeasure: Employee training with standards of business conduct
Risk 18: After the banner ad designs have been approved, a small number of them (3%) are corrupted and nobody notices. If the corruption is just garbage data, then Nullco is paying for banner ads which will not be effective. If the corruption is replacing the images with offensive content, then Nullco has a reputation problem as well. Countermeasure: process control with data quality validation.
Risk 19: A disgruntled or bribed/blackmailed employee at one of Nullco's service providers reveals sensitive information from the SuperWun repository to other companies using the access right that was granted by Nullco, or posts these details to the internet or a news site, or deletes or alters data. Countermeasure: Watermarking of data to deter data leaking

3.4 Selected Threats to asset: Mercado Portal

Threats from accidental causes
Risk 20: People share a large volume of documents and information when collaborating using this portal. It is possible to share the wrong document by mistake and expose information to those who shouldn't see the information. For example, someone negotiating a contract for services may be negotiating with several parties for the same service trying to get the best deal. They could send proposed terms being worked with provider A to provider B and damage Nullco's negotiating position with them. Countermeasure: Automatic tagging of documents with metadata and access control policy enforcement using this metadata (which limits to whom the document can be sent).
Threats from deliberate causes
Risk 21: A person in Nullco is authorized to select services using the Mercado portal and commit to contracts. Nullco has a policy which the Mercado portal enforces that each person may only sign contracts for up to $10,000. For contracts from $10,000 to $100,000 two people need to sign, and above $100,000 three people need to sign. An employee could get around this limit and bypass Nullco's policy by signing several contracts each worth $10,000, all on the same day or the same week, because the employee is aware that Nullco cannot enforce a limit on the total dollar value of contracts per time period that a person can authorize. An employee who will soon take another job at a different company may violate this policy without being caught and then leave Nullco. Countermeasures: Continous monitoring and process control
Risk 22: Nullco employees should select services based upon what will be best for the company overall. However, someone could select services based upon what will be easiest for them or their department without regard for the overall impact to the company. This causes inefficiencies and lower profits, partially because Nullco does not have a solution of forcing the decision maker to document their rational for making the decision, based on the criteria that is designed in the service selection process. Countermeasure: process control to force documenting the rationale, which is required as part of the service selection process.
Risk 23: The Mercado portal provides information on customer reviews of the services hosted in Mercado. These reviews become publicly visible. This provides valuable information to Nullco's competitors on who Nullco's satisfied and dissatisfied customers are so that they can poach Nullco's customers and impact Nullco's business. Countermeasure: A well-defined process to review and declassify potentially sensitive customer data

3.5 Selected Threats to asset: data held at SuperAcim service

Threats from accidental causes
Risk 24: A SuperAcim employee incorrectly configures the configuration table for the database such that it points to a data source which holds information for another company instead of Nullco's when building the data model. This results in a poor quality model. This model is sent to another department which uses it on the right data source to generate the campaign management rules. Because the data model is wrong, the business rules are wrong, and the campaign will not be very effective. Nullco does not have a way of detecting this error until late in the campaign. Countermeasure: information flow control on reading/writing sensitive data
Threats from deliberate causes
Risk 25: An employee has a spouse that works for one of Nullco's channel partners, and gives the SuperAcim report on Nullco's supply chain management related data to his spouse to help the spouse's

career, which causes Nullco to lose sales to this partner. A similar incident could happen if a competitor convinces one of Nullco's employees to dupe the co-worker that can access the sensitive information to reveal information. Countermeasure: Employee training on standards of business conduct
Risk 26: SuperAcim could purposely give Nullco poor quality data since Nullco competes with them in some market segments. For example, they remove some of their very valuable customers from the lists they send to Nullco. Countermeasure: continous monitoring of database changes at SuperAcim with auditing to detect such deletions

4 Service Customer Requirements

In this section, we categorize the countermeasures proposed for risks identified for the Nullco Product Marketing Campaign and examine their implications on shaping the service customer requirements. In addition, we follow the expected lifecycle of service customers' experience in the shared marketplace, to evaluate how and when these service customer requirements need to be satisfied by the service providers.

4.1 Detailed Customer Requirements

We categorized the risks and countermeasures to better understand customer requirements. The categorization is shown in Table 2.

Table 2: Categorization of Countermeasures for Selected Risks

Countermeasure Category	Risks Addressed
Quality Assurance of Services & Data	Risk #5, 7, 8, 12
Proper Usage Control of Data	Risk #1, 19, 20, 24, 26
Sufficient Business Process Control	Risk #3, 4, 6, 8, 13, 14, 16, 18, 21, 22
Trust through Transparency	Risk #9, 26
Well Established Risk Management	Risk #2, 12

A services marketplace can increase the efficiency of participating companies by allowing them to focus on what they do best and outsource non-core aspects of their business. In the existing service outsourcing paradigm, an outsourcing service provider establishes their reputation and trust with their customers over a long period of service engagement, and by following mature IT governance standards such as COBIT [4]. In contrast, the new and dynamic services marketplace hypothesized in the Mercado project poses a myriad of challenges to participating companies, because customers may have less control and transparency of processes, and little or no experience with a potential service provider. Customers will have needs which are difficult to quantify and meet:

- Adequate Quality Assurance of Services and Data: The customer needs to know that the quality of both the services and data that they receive is adequate for the purposes for which they are used.
- Proper Usage Control of Data: Customers need to trust that their data is safe being used by many providers. How can they be sure that the data controls expressed in the contract are being followed

accordingly when many different service providers are likely touching the same piece of data in order to provide the end-to-end service capability to the customer? They need recourse if their data is compromised.

- Sufficient Business Process Control: Customers don't want a mistake or malicious act by their own employees or a service provider's employees to cause damage to their business. They need to be assured that adequate process control is in place. For example, the customer needs to ensure that there is a proper process control to address the issue that an error or malicious act by an employee doesn't commit them to a contract to which they don't want to be committed.
- Trust through Transparency: By adequately revealing to the customer the continuous running of their business activities behind the service interface, the marketplace can build trust in the services in a better way because customer concerns can be resolved through a transparent process, which is better than a binary answer from a third-party certifier.
- Well Established Risk Management: Customers want to prevent problems from occurring whenever feasible. When problems do occur in case all the measures fail to prevent them, they need to quickly detect the problems with the data or services *before* they cause serious problems for the business (including reputation damage, loss of customers, lawsuits and penalties). There is opportunity to find techniques and build systems that help avoid bad consequences resulting from human or machine error. By detecting the error early, it can be quickly corrected before it snowballs into a much larger problem. It is likely that the same techniques can be used to quickly detect and reduce the impact of malicious activity as well.

Furthermore, there is a requirement to provide the customers with **assistance on business practices in the marketplace.** The customer will be unfamiliar with a services marketplace, and will need assistance to get a business process running using this new paradigm. They require assistance at two levels. Firstly, they require help to select and compose services which provide the desired capability. Best practices and reusable patterns can be captured in templates to make this task easier for first-time customers. Secondly, they may require help to customize selected services to find optimal (or at least adequate) service configuration options for their needs.

4.2 The Customer Engagement Model

Customer requirements become more rigorous as their business becomes more dependent on services from the marketplace. We classify the customer engagement with a services marketplace into four time slices in order to determine what is necessary to gain the customer and move the customer to become an active long-term user of the service:

1. **The trial period.** Step one is to get the customer to try the service. A customer will usually use a non-critical business process to test the service during the trial period. They may even use fake business data for the first trial. But eventually they will need to try the service using real business data to ensure the service can meet their needs, and therefore this requires that the service provider deliver the first two needs above, Adequate Quality of Services and Data, and Proper Usage Control of Data. Free trial periods can help entice the customer to try the service, but the cost to try a new and unproven service is measured in more than dollars, due to staff time, opportunity cost, business risk, and other factors. In order to convince a customer to begin using a service, a satisfied reference customer can help immensely. The reputation of the service provider will also play a role.
2. **The continuation period.** When the customer tries the service, it must be an adequate and cost-effective solution to the business process or problem at hand in order for the customer to continue using the service instead of reverting to their prior solution. What we call the *visible attributes* of the

services must satisfy the customer and improve upon their prior solution. Visible attributes are things the customer can easily verify for themselves, like acceptable functionality, performance, support, and cost in delivering the business process. During this period, it is important for the service provider to deliver the third and fourth needs above, Sufficient Business Process Control and Trust through Transparency.

3. The expansion period. In order to expand usage of services, and in particular to make the success of the business dependent on services, customers need to feel confident that all their needs will be met. Good past experiences and trust through transparency will help, but will not be sufficient to address this need. The final need identified above, Well Established Risk Management, will be necessary for a customer to take this step. The customer will need to gain confidence that what we call the *invisible attributes* of the services are adequate for their needs. Invisible attributes are things that are difficult or impossible for the customer to verify for themselves, but nevertheless are important to the customer. This can include things like the long-term availability and security characteristics of the service. The customer can sometimes observe that certain invisible attributes are inadequate due to failures such as downtime or security breaches, but often the customer doesn't have the expertise or knowledge to verify that the invisible attributes are adequate for the needs of their business. Thus, there is an opportunity to provide various verification services which vouch for the sufficiency of the invisible attributes of other services for a particular use. For example, a verification service could warn the customer of a potential availability problem to avoid Risk 12.
4. The champion period. Some customers will like the service so much that they become fervent promoters of the service, referring their colleagues and partners to the service. Services can take advantage of this by enabling customers to publicly tell their stories of success while using the service, which provides mutual benefit. Further, such satisfied customers can share their experiences with the newcomers to effectively expand the community of the services marketplace.

Figure 2 shows the mapping of the four customer engagement phases described above with the identified service customer requirements.

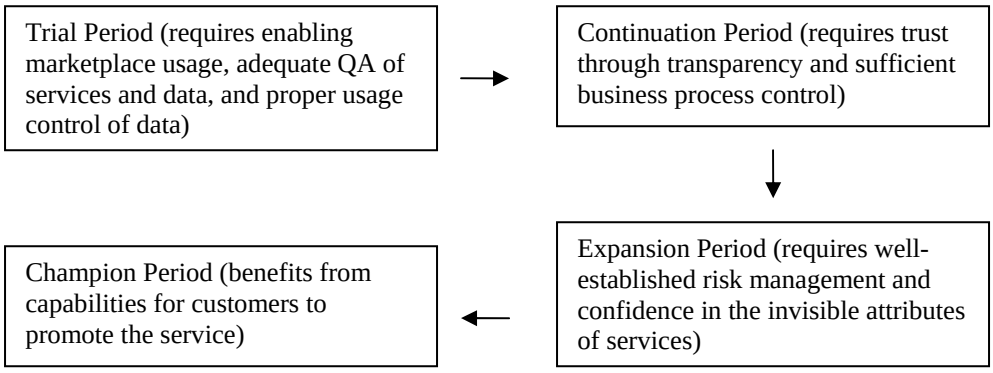

Figure 2: Mapping of customer engagement phases to customer requirements

5 Identified Research Areas and Context

From the threat analysis presented in Section 3 and customer requirements identified in Section 4, we have identified three research areas which we detail in Section 5.1. We position these research areas in the environmental context of the services marketplace in Section 5.2.

5.1 Research Areas

Business Data Management. A fundamental assumption of Mercado and similar dynamic service portals is that a service can be substituted by other services equipped with the same or greater functionality. The core business asset that needs to be protected to ensure the continuity of the running business is the data that is created, processed, and maintained by the services. A broad spectrum of techniques is required to correctly manage business data. Data access control serves a proactive purpose and it requires data access policies at finer granularity by following the principle of least privilege (Risks 3 and 15). Aside from access control, the issues of data appropriateness checking (Risks 3, 4, 5), data quality checking (Risks 7, 8), still remain. Finally, information leak prevention and detection (Risks 19, 23) can help ensure that the data being accessed does not fall into the wrong hands and should a leak occur the culprits can be identified.

Continuous Business Monitoring & Validation. To improve the confidence of the service owner and the trust by the service partners and customers, each service would need to have transparency by having its running business activities monitored, either by itself through a white-box monitoring mechanism, or by third-party companies through a black-box monitoring mechanism (Risks 4, 9, and 26). A set of rules, defined either by a single organization, or by different organizations with contractual business relationships, needs to be validated to ensure that the running business follows the intended design and provides the intended capabilities with the necessary quality. These rules deal with business contractual relationships (such as service agreement), fraud detection (in service outsourcing), information leaking (Risks 19 and 23), and regulatory compliance. As shown in Risk 3, data management policies can be potentially cleverly bypassed by insiders. Thus active monitoring serves a complementary purpose to detect policy violations.

Process Control on Data Management and Continuous Monitoring. The aforementioned two research areas lead to various point technologies that when combined, could be more powerful. For example, information leak detection, continuous monitoring, and quality assurance can be combined with access control, as shown in the countermeasure for Risks 3 and 15. To bridge the IT-business miscommunication gap, a business process or process control needs to be linked to high-level business objectives such as security, regulatory compliance and business assurance at the business specification level, and correspond to a concrete enforcement mechanism at the system implementation level, which conducts runtime data management and continuous monitoring. Finally, process controls should be captured in templates for individual services and customized for each service customer due to the customer's operational constraints and preferences.

5.2 Context of the Services Marketplace

Similar issues and research problems have been identified and well explored in a service computing environment, albeit with focus mainly within a single enterprise domain. The services marketplace imposes a host of new technical challenges on business assurance, due to the following unique environmental constraints.

Cross Organizational and Decentralized. A business process in the services marketplace can be composed of many services provided by many organizations. Policies defined in one administrative domain to manage the service and data access within that domain, should be applicable to the service access requestors that come from a different organization in a different administrative domain. Centralized policy enforcement frameworks, such as the ones that are based on federated identity management, can be unmanageable and un-scalable, given the sheer number of different companies

hosting services in the services marketplace. Therefore a decentralized approach is required in this new cross-organizational environment to conduct data management and continuous monitoring. To support such a decentralized approach, access rights management will encounter issues like how to grant the rights, how to manage and delegate the rights, and how to enforce the rights. In the case of continuous business activity monitoring, events can come from multiple organizations. The policies to govern the publication/subscription of the events as well as the reception of event notifications will be dependent on individual organizations and also the inter-organizational business contractual relationships.

Dynamic Service Composition. The same or similar service functionality can be provided by alternative service providers and therefore theoretically a service can be replaced dynamically due to cost, performance, or reputation, but the business application can still meet its requirements. In a changing environment, data management policies might be changed due to the structure of service composition. For example, the content update policy may require approval from three organizations: the owner of the content management service, the service consumer, and the service that provides content creation. When the content creation service is replaced, the approval process needs to be updated as well. Similarly, the rights of continuous business monitoring-related event publication, subscription, and notification will depend on which target service providers the business has chosen, due to different levels of trust and inter-organizational business contracts.

Service Mediation and Sub-contracting. Businesses like long-term relationships, and the trust they build up over time can be more important than getting the best price each time a service is used. Mediators offer such capabilities. The mediator assumes the risk from the mediated services, and the customer gets the benefits of a much more dynamic market but still has the single point of contact for support, billing, dispute resolution, and other services. Subcontractors can dynamically attach and detach their services to the mediator service. Sub-contracting can introduce many issues. Access rights inheritance from the mediator is subtle, for example, when dealing with inter-country data migration [5]. A particularly vexing issue is the question of terminating long-lived access rights. Services often are long-lived objects but the actual nature of interactions between the consumer, service provider, and sub-contractors changes over the lifetime of the contract. It is important that sub-contractors have access to exactly the data and control that is needed for their function at any point but this is hard to describe or define, let alone enforce. Thus, granting access to data becomes much more difficult than revoking that granted access. A continuous monitoring framework will likely be reconfigured dynamically due to the change of service composition structure. An as yet unresolved problem arises when there is an inherent conflict between customers who cannot share sensitive data with the service provider without violating their internal privacy or business competition policies. In some cases a service provider who deals with a client's very sensitive data may not be allowed to subcontract components of their service.

Multi-tenancy and Packaged as Software as a Service (SaaS). In a services marketplace, service functionality regarding data management, process control management and continuous business monitoring and validation can potentially be packaged as services, similar to authentication, authorization, and even vulnerability management [6]. In the multi-tenancy environment [7], each consumer has its own view of data confidentiality, data integrity and data availability, and thus its own unique management policies, methods of packaging the policies into business workflows, and objectives of running the business and thus a completely different way of monitoring and validating business activities. All these issues dramatically lead to management complexity. Furthermore, corporate firewall traversal can prevent bi-directional data exchange between the behind-the-firewall client and the open Internet service, especially when transferring sensitive and private data [5].

6 Related Work

The research areas defined in Section 5 cover a broad spectrum of engineering disciplines. In this section we survey the research and development efforts in the context of decentralized service-oriented environments that involve multiple organizations.

Federated Identity Management (FdIM) [8] supports cross-organizational single-sign-on. It addresses only the cross-organizational authentication issues. The authorization decision still relies on the policy databases hosted in individual organizations and the requestor being authenticated locally by these organizations. Various access control policy frameworks are available to handle cross-organizational access. Most of them are identity-based, with variants being role-based and attribute-based [9, 10, 11]. These identity-based mechanisms work well with FdIM but suffer various identity management related problems, such as granularity of right permission, ambiguity of right enforcement, right delegation, etc. The authorization-based access control offers an alternative mechanism to address these issues [12].

In terms of database outsourcing, various techniques have been developed, including data encryption on databases [13], data integrity validation by inserting small amounts of pre-computed records into databases [14], and data privacy preservation for aggregation-based queries based on statistical perturbation of query results among multiple database providers [15].

A model-based assurance framework [16, 17] captures the IT controls as models. Each model identifies the data sources that represent the outcome of the modeled IT controls, provides analysis routines on the identified data sources, and defines analysis result presentation formats. The framework allows certain auditing processes to be automated once the controls are captured, and thus facilitates continuous controls monitoring. Similar continuous monitoring facilitates corporate auditing by analyzing financial transactions (payroll, purchase-to-payment, etc.) continuously [18]. Currently such techniques are mostly concerned with a single enterprise's internal auditing. How to incorporate multiple services hosted by different providers to achieve cross-organizational business assurance still remains open.

Platform support for SaaS has been provided by Salesforce.com, Amazon.com, SAP, Microsoft, IBM, Google, etc. Multi-tenancy support to facilitate customization of data modeling, business workflow orchestration, and even end-user application UIs, plays a key role [7, 19, 20]. However, all these platforms are positioned mostly for single hosting service providers, even though the services hosted in their environment can call out to or be called by services hosted elsewhere. Dynamic service composition requires distributed and cooperative controls over data to achieve business assurance, and thus a new set of programming abstractions and runtime support is needed. The Apex platform [20] relies on database role-based access control because the applications are inherently web-enabled database applications. IBM advocated their identity-management support in the SaaS environment [19], with the belief that what is offered for a physical enterprise can also serve well for a virtual enterprise.

7 Conclusions and Future Work

A services marketplace as proposed by Mercado can potentially revolutionize the development and operation of business applications over the Internet. But in order to make it a viable place where customers are confident doing business, many challenges need to be tackled beyond providing a scalable runtime infrastructure with a user-friendly development environment. A key challenge is providing business assurance—how to ensure the customers that they can trust the marketplace and the services offered in the marketplace to store, process, transfer, and present their data in an appropriate manner aligned with their individual business needs. This report illustrates this challenge by conducting a

threat analysis on a fictional company called Nullco conducting a product marketing campaign, with its reliance on multiple services selected from the services marketplace.

In this report, we highlighted possible countermeasures from the identified risks, categorized them, examined them, and then derived the customer requirements based on the identified risks and countermeasures. We then selected and described three research areas: business data management, continuous business monitoring and validation, and process control for business data management and continuous business monitoring/validation. In particular, process control is critical to bridge the current IT-business miscommunication gap by linking high-level business requirements with low-level concrete enforcement mechanisms.

Following the requirements that we have identified in this report, we have developed solutions which address some issues in these research areas. A database fingerprinting methodology [21] can detect leaks of database information by service providers based upon the insertion of artificial records in databases. A unified resource lifecycle management framework [22] allows customized data management policies to be enforced by customizing the policy templates defined by service providers. Apart from continued improvement and prototyping of these solutions, we believe that the area of continuous business monitoring and validation still remains open, but is key to providing business assurance for service partners and customers.

Our overall objectives are to improve customer confidence when using services over the Internet and enable easier and more automated use of such services. We believe that there will be an abundant set of other technical problems and solutions unique to this services marketplace environment.

8 Acknowledgements

We wish to thank the team which created the use case which we used: Don Young, Jim Pruyne, Kivanc Ozonat, Martin Arlitt, Sujoy Basu, and Sven Graupner. The authors are grateful to Raj Rajagopalan and Martin Arlitt for reviewing and providing many suggestions for improvements to this report.

9 References

1. J. Scheck, “HP Sees Everything as a Service”, http://blogs.wsj.com/biztech/2008/05/20/h-p-sees-everything-as-a-service/.
2. Octave-S, http://www.cert.org/octave/osig.html.
3. B. Stephenson and J. Li, “Product campaign threat analysis,” http://opra.hpl.hp.com/Mercado/MercadoTA04.doc and http://opra.hpl.hp.com/Mercado/MercadoTA07.doc.
4. ITGI, Control Objectives for Information and Related Technologies (COBIT), Fourth Edition, 2005.
5. J. A. Mulligan, “Best Practices: Managing Global Privacy Programs,” Forrester Report, Sept. 2007.
6. Qualys, http://www.qualys.com/.
7. F. Chong, G. Carraro, and R. Wolter, “Multi-Tenant Data Architecture,” 2006, http://msdn2.microsoft.com/en-us/architecture/aa479086.aspx.
8. Liberty Alliance, http://www.projectliberty.org/.
9. M. Blaze, J. Feigenbaum, and J. Lacy, “Decentralized trust management,” Proceedings of IEEE Symposium on Security and Privacy, pp. 164-173, 1996.
10. M. R. Thompson, A. Essiari, and S. Mudumbai, “Certificate-Based Authorization Policy in a PKI Environment,” ACM Trans. Information and System Security, Vol. 6, No. 4, Nov. 2003.
11. J. Joshi, R. Bhatti, E. Bertino, and A. Ghafoor, “Access-Control Language for Multidomain Environments,” IEEE Internet Computing, Vol. 8, Issue 6, Nov.-Dec. 2004.
12. J. Li and A. Karp, “Access control for the services oriented architecture,” Proceedings of the 2007 ACM workshop on Secure web services, pp. 9-17, 2007.
13. H. Hacigumus, B. Iyer and S. Mehrotra, “Providing database as a service,” In ICDE, pp. 29-38, 2002.
14. M. Xie, H. Wang, J. Yin, and X. Meng, “Integrity auditing of outsourced data,” Proceeding of VLDB’07, pp. 782-793.

15. L. Xiong, S. Chitti and L. Liu, “Preserving data privacy in outsourcing data aggregation services,” ACM Trans. On Internet Tech., Vol. 7, No. 3, Article 17, Aug. 2007.
16. A. Baldwin, Y. Beres, S. Shiu, “Using Assurance Models in IT Audit Engagements,” HPL Technical Report, No. HPL-2006-148R1, 2006.
17. 17. A. Baldwin, M. Casassa Mont, Y. Beres, and S. Shiu, “On Identity Assurance in the Presence of Federated Identity Management Systems,” Proceedings of the 2007 ACM workshop on Digital Identity Management.
18. ACL, http://www.acl.com/.
19. IBM Software as a Service, http://www-304.ibm.com/jct09002c/isv/marketing/saas/
20. C. Roth, D. Carroll, and N. Tran, Creating On-Demand Applications: An Introduction to the Apex Platform, free online electronic book accessible via Salesforce.com.
21. E. Uzun and B. Stephenson, “Security of Relational Databases in Business Outsourcing,” HP Labs Technical Report HPL-2008-168.
22. J. Li, B. Stephenson, and S. Singhal, “A Policy Framework for Data Management in Services Marketplaces,” submitted for publication.