

2009



Hewlett-Packard
Point of View

Mark Jacobsen, Mick
Keyes & Christian
Verstraete

[BEYOND FOOD RECALL, THE BUILDING OF A CLOUD COMPUTING PLATFORM]

This document describes how we started from the requirements for food recall to develop a Cloud Computing Platform for the Manufacturing Industry. After describing the specific requirements of recall, we explain how those got implemented and how the developed platform could help address different but related problems.



BEYOND FOOD RECALL, THE BUILDING OF A CLOUD COMPUTING PLATFORM

MARK JACOBSEN, MICK KEYES & CHRISTIAN VERSTRAETE

Over the last couple years, several food recall initiatives have hit the press and gotten major exposure. Unfortunately a number of these have ended with casualties and huge costs for the suppliers. It became clear there was a need for a more efficient process of exchanging data across the food ecosystem and managing food recall. Several pilots have taken place over the last 24 months, but most have demonstrated a number of limitations making them impractical for being rolled-out worldwide. The “cost of pain” for the US food industry is estimated¹ at:

Beverage Industry Product Recall Cost (estimate)	\$250 Million/Year
US Societal Annual Costs from Food Borne Illness	1-100 Hospitalizations, 1-500 Deaths \$6.9Billion/Year Costs
Recent "tomato" Salmonella Case – CDC states Probable Cause : <i>Salsa, Cilantro or Peppers</i>	Florida tomato Ind. est. Cost - \$40 M California & Mexico ind. Cost expected to reach > \$100M
Global Food/Beverage Industry Counterfeiting Costs	\$60 Billion/Year
Global Food/Beverage Industry - - Brand Protection Market Value Forecast	2006 = \$2.9 Billion 2012 = \$5.0 Billion

HP got approached to help addressing the issue by developing an environment that focuses on managing flows of goods through the supply chain and facilitates the management of a recall process. HP had to take into account amongst others the latest legislation in the US, known as HR2749², currently being finalized.

¹ Source : US CPSC, US FDA, USDA, US EPA, Global insights (IPG Research), US ERS, Pira International

² http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=111_cong_bills&docid=f:h2749fs.txt.pdf

THE REQUIREMENTS

Resulting from increase of food recalls and scandals around the globe over the last couple years, the need for a careful management of food across the supply chain, from farmer to consumer, is becoming increasingly critical to provide a fast and focused response. As the US proposed legislation points out, a tracing system enabling *“to identify each person who grows, produces, manufactures, processes, packs, transports, holds or sells food in as short of a timeframe as practical but no longer than 2 business days”* is required. This results in the need to:

- Maintain the full pedigree of the origin and previous distribution history of the food
- Link that history with the subsequent distribution of the food
- Establish and maintain a system for tracing the food that is interoperable with the systems established and maintained by other such persons
- Use a unique identifier for each facility owned or operated by such person for such purpose

These needs imply the existing of three basic elements:

- The establishment and maintenance of lot numbers
- A standardized format for pedigree information
- The use of a common nomenclature for food.

Multiple approaches of this problem are possible and the first one coming to mind is the development of a food industry hub to centralize all information related to the movements of food. However, discussions with governments around the globe made very clear the willingness of many of them to want to keep their food information under their own control, making the centralization of that information impossible. As an increasing amount of food crosses several boundaries before being consumed, it became quickly clear to us that we needed to develop an environment capable of coping with distributed data. For this reasons we looked at a “cloud computing³” based approach to address the needs of the food industry. In line with the NIST, Information Technology Lab definition, referred to above, we actually created a *community cloud*⁴, as it allowed us not only to address the legal requirements established above, but also to take into account the need for distributed data. The way the system was architected gives us the opportunity to develop other services, with similar requirements, moving forward. We also teamed up with GS1⁵ to ensure the common nomenclature; the unique facility identifier and the interoperability with most of the existing systems as GS1 have played an important standardization role in the food industry.

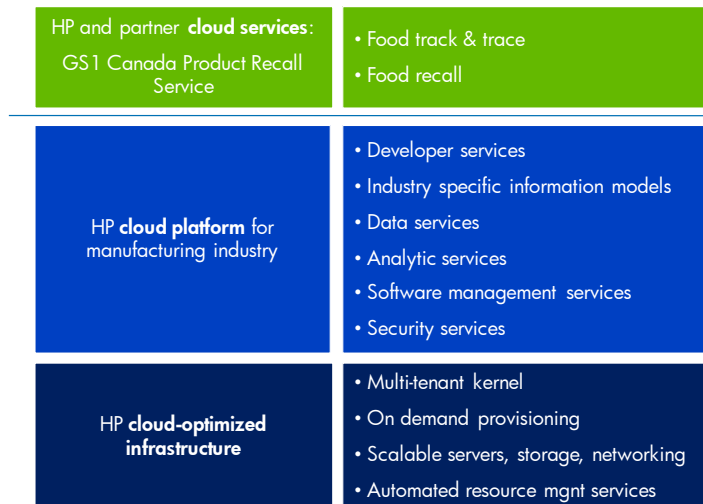
³ <http://csrc.nist.gov/groups/SNS/cloud-computing/>

⁴ The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premise or off premise.

⁵ <http://www.gs1.org/>

THE OVERALL ARCHITECTURE

To address the requirements of the previous page, we developed a three layer architecture highlighted in the picture. As



stated earlier, this architecture allows the development of a number of services. We plan to take advantage of this in the future.

The first layer, we call the “cloud-optimized infrastructure” includes beside the physical hardware on which the services will run, the provisioning, multi-tenant, and management capabilities. This is critical to develop the sound foundation for a reliable service.

The second layer, called “cloud platform for the manufacturing” contains the basic services that are required for the development of the composite service that is actually presented to the community. Beside generic services such as security & software management, services more specifically geared towards the collaboration of partners across an ecosystem have been implemented. Such services include data services, allowing the access of distributed data sources, and analytic services.

The third layer is the “cloud services” layer and delivers the specific service(s). Obviously, there can be multiple services provided using the same underlying layers, as long as the services have similar requirements from an infrastructure and platform perspective. By separating the architecture in the way described above, we have actually built an environment allowing us to develop new services quickly and an environment that can be operated at low cost. The developments were made by our Cloud Computing center in Galway, Ireland, drawing on research performed by HPLabs.

The cloud services layer interacts with the user’s browser, providing access and navigation capabilities, and manages the food recall process in this particular instance. It draws on the generic services available in the underlying layers to deliver the functionality required.

In the next section we are describing the infrastructure and platform layer in more details and in the subsequent section we will discuss more specifically the data aspects as these address critical requirements of the Food Recall service.

DESCRIPTION OF THE UNDERLYING LAYERS

THE INFRASTRUCTURE LAYER

The infrastructure layer is based on HP's Adaptive Infrastructure Services (AIS) offering, an IaaS offering, built using the latest HP BladeSystem c-Class and Integrity servers in various standard configurations including virtualized O/S instantiations. Although AIS offers a choice of operating systems, we choose to use Microsoft Windows to build the service described here. The AIS Data Center network fabric includes a high-capacity switching fabric, routers, firewall devices and VPN termination equipment. The latter feature allows the users of the service to interact using secured VPN connections. Access controls are policy based and implemented with a combination of dedicated VLAN's, Access Control Lists, packet filters and stateful inspection firewall devices. Management tools are housed in a dedicated network compartment which is accessed by HP support personnel via secure shell (SSH) for UNIX or Linux or Remote Desktop for Windows. Authentication to this management network is two factors and access is logged, increasing the protection of the environment. Storage Area Networks and data protection components are responsible for providing fiber channel based storage, back-up services and the underlying fiber channel network utilized by the servers and applications. The environment is preconfigured with redundant power, network and storage connections within the data center, and a back-up environment takes over in case of fail-over of the initial one.

THE PLATFORM LAYER

The platform layer used to build and launch the Recall service is based on shared foundation service layer and a Service delivery layer which are the basis of the HP Cloud platform for manufacturing offering.

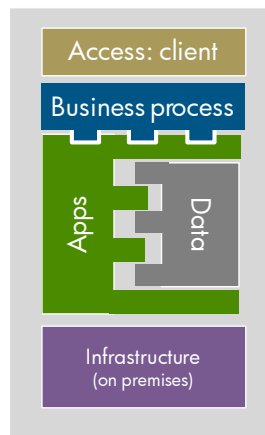
The platform layer supports multiple development and lifecycle management. The HP SOA centre suite encompassing HP BTO (Business Technology Optimization) software is central to the platform. It provides complete service lifecycle governance capabilities, including an authoritative system of record, customizable information management, automated lifecycle management end-to-end policy management and consumer/contract. The platform is provisioned to launch specific services. For launching the specific Recall service we used a .Net development framework. For data access an EPCIS compliant module based on Microsoft BizTalk was used as the application integration component with SQL as Database. BTO software components were used for management.

The overall Cloud manufacturing platform layer also includes industry specific modules in the area of Visualization which offers a variety of services such as Visual Business intelligence for supply chain. It will also include a Data fusion layer which offers the ability to define and launch powerful services from distributed data stores

HANDLING DISTRIBUTED DATA

One of the key requirements of the Food Recall service is the need for being able to draw upon distributed data sources. We believe other business problems would benefit from a similar approach. These include areas such as supply chain collaboration, hazardous material reporting and others. In this section we will describe how a cloud based approach allows us to address such requirements, but let's first look at more traditional approaches and what issues are related with them.

A TRADITIONAL BUSINESS PROCESS APPLICATION




Most applications used in manufacturing today have been developed using traditional “client-server” approaches and expect to be run “on premises”. Data is maintained in a database that runs on the same or a separate system and that directly interacts with the application, typically through the use of SQL code. The business processes in turn are programmed, or in the best case, configured in the application. A presentation layer (not highlighted in the diagram) handles the interaction with the user. This works very well for applications whose scope is limited to the boundaries of a single enterprise and whose workloads are stable. In many situations, dedicated servers are reserved for specific applications.

As enterprises have increased the outsourcing of some of their activities, integration of information across company boundaries have become more important. During the internet boom of the early years of the millennium, many start-ups have proposed a “hub” based approach. This one consisted fundamentally in recreating a dedicated environment (within the premises of the start-up), increasing the levels of security to allow multiple companies to access the same infrastructure and functionality. However, it required the data to be located within the “hub”, forcing the users to duplicate their information. This in turn resulted in information being out of sync, in not really controlled proliferation of some of that data and in potential security breaches. As multiple companies were addressing the same space, it became way more difficult to predict the capacity needed in the hub. Most of these start-ups have disappeared in the internet meltdown, as they never could get their business model right, but some remain active to date.

STANDARDIZATION & VIRTUALIZATION

The unpredictability of the workload can be addressed through standardization and virtualization as long as the application has been developed to take advantage of these approaches. Virtualization addresses the performance aspects, but does not address the data issues in any way.

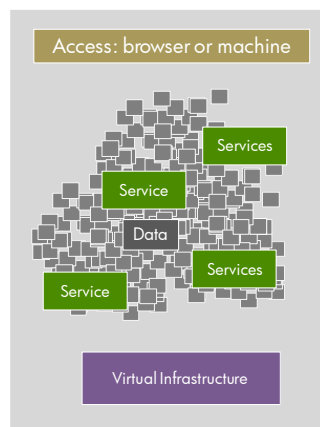


When talking about cloud, it is this standardization and virtualization however that is highlighted. The possibility to ramp up and ramp down servers quickly (called elasticity) is put forward and examples are quoted on how start-ups have been able to expand from 50 to 5000 servers in a matter of hours. However, this is not the typical need of manufacturing companies. Yes, the ecosystem workload needs are varying, but the variability is reasonably small in the collaboration space. The issue is data, and virtualization & standardization still assume a single database interacting with application programs and business processes. The programs are now

developed in such a way they are parallelized over multiple servers, and from that perspective, they have been optimized, but that's it for the changes.

DEVELOPING A BUSINESS ECOSYSTEM

If data is kept in a single place, the most up-to-date information is always used. If data is kept under the responsibility of the owner, sharing the information becomes more acceptable as the access security is managed by the source, which is under control of the data owner.



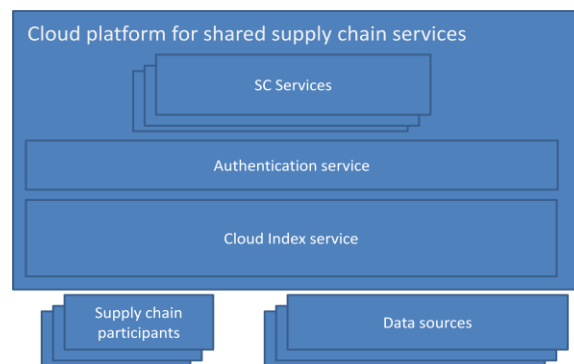
As companies thoroughly look at how reduce costs in their ecosystems, by reducing safety stocks and buffer inventories, and as such making the supply chain leaner, they are removing the elements shielding them from issues higher up in the chain. As their customers still require the same predictability of delivery, they are forced to find other ways to address disruptions. Sharing information as soon as something happens, allowing faster reaction is one of the ways to address these requirements. But

suppliers are often hesitant to share information by transferring data to hubs as described above.

Having direct access to the data source from services running in the cloud could address these problems and fears. This is an important, but often overlooked, aspect of cloud computing, which we have used as part of the cloud computing platform we developed for the food recall service.

ACCESSING AND MANAGING DISTRIBUTED DATA IN THE CLOUD

In traditional environments such as B2B Hubs, data is uploaded in the hub and made available to all participants. We call this a push model. Unfortunately, for a number of reasons described earlier in this document, such approach is not applicable in the current environment as information ownership and governmental regulation amongst others require the data to stay within particular boundaries. Delivering this service worldwide required us to look at how we could centralize the appropriate data on a need to know basis, which we call a pull model.



In a pull model, the data remains under the responsibility of the initiator or owner, and is only transferred to the cloud based service at request. Data resides in the cloud only for the purpose of one service transaction and is then whipped out (Transient Data). This requires us to address two particular issues, first

locating the data and second, extract the data from its original source and put it in a format that the cloud service understands. To achieve this we have developed a Cloud Index and Data Discovery Engine, often referred to as “Cloud Index Service” or CIS. Once authentication has been performed, CIS identifies where the data is located, and in what format it is provided. CIS is also able to extract the data from the source in a format useful to the service. All services built for the Supply Chain have to depend on the Cloud Index Service for their data. They are not allowed direct access to the data sources.

As part of the partner on boarding process to the cloud service, the partner will identify which data is made visible to the service, in what format that data is and where it is located. This serves as a base for the index. At that moment, the owner may decide to provide the service with direct access to his data marts, identifying the public and private data items, or he could decide to set-up a shadow database in which he only puts the data available for the service. In doing so he would secure his private information. This is totally transparent to the service.

To build trust between the owners of the data and the service providers, the workings of the CIS need to be totally transparent and open to the supplier of the data. Once a query is performed in the cloud, an agent is generated (eg. A Java Applet) to query the owner databases for the relevant data, as only that data should be transferred centrally for the purpose of the overall query. The transfer should be transient, which means that, once the report has been generated, the data is destroyed. The agent performs the query on the partner systems directly and packages the data in a format that is under-

standable by the cloud environment. In the product recall service we have used the EP-CIS standard as the format, but we could imagine working with XML standards or others. However, for this to work a unique identification of the parts needs to exist and all partners have to include that in the information they maintain. The agent will then be able to send the relevant data in the cloud in a standardized format to be used by the analysis service.

Moving forward, we are looking at providing services integrating both the push and the pull model. This may allow smaller players to host their information in the cloud, and free them from the obligation to have to maintain their own IT environment. The CIS will be expanded to access both the cloud and distributed data.

This model allows the services to be written using a defined data access model, regardless of the data sources they are accessing. It also allows for a form of authentication to be built in controlling access to the data sources.

Supply chain members once they decide to collaborate, can without any massive outlay of hardware and infrastructure provide access to the data they wish to share and then select which services they want.

CONCLUSION

To address inter enterprise collaboration, manufacturing companies need to exchange information and have access to each-other's data. As described in this paper, the cloud provides mechanisms to maintain the data in the control of the originator while providing secure access to the ones needing to know. This feature provides a unique opportunity to develop a new set of services geared at improving collaboration between enterprises in an ecosystem.