IBM Research Report

Taking IT Management Services to a Cloud

Michael R. Head

Grid Computing Research Laboratory
Binghamton University
Binghamton, NY 13902

Anca Sailer, Hidayatullah Shaikh

IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Mahesh Viswanathan

IBM Global Technology Services
294 Route 100
Somers, NY 10589

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

Taking IT Management Services to a Cloud

Michael R. Head[1],   Anca Sailer[2],   Hidayatullah Shaikh[2],   Mahesh Viswanathan[3]

[1]Grid Computing Research Lab,
Computer Science,
Binghamton University,
Binghamton, NY 13902 USA
mike@cs.binghamton.edu

[2]IBM TJ Watson Research Center,
Hawthorne Research Lab,
19 Skyline Drive,
Hawthorne, NY 10532, USA
{ancas,hshaikh}@us.ibm.com

[3]IBM Corporation
Global Technology Services,
294 Route 100,
Somers, NY 10589, USA
maheshv@us.ibm.com

Abstract

While IT management services represent a mature subject in the IT business arena, the emerging cloud generation of management services require critical enhancements to the current processes and technologies in order to deliver IT management remotely with rapid on-boarding and minimal labor involvement from experts, to be affordable and scale up to the promise of the cloud. Traditional Remote Infrastructure Management (RIM) service providers use their own Network Operations Centers (NOC) to remotely monitor and manage customers' IT infrastructure. The primary business value for RIM services is that it helps global enterprises to small and medium businesses (SMB) to outsource the burden of managing their IT infrastructure, thus cutting down costs for infrastructure management and gaining access to expert skills. Although the IT management service itself delivered this way is more affordable, the RIM customer on-boarding process particularly is not, taking between one to two months of expensive labor. The multiple face to face or phone interactions with the customer to discover their IT environment, identify the resources to be managed and guide the enablement of the environment for remote management is proving increasingly ineffective and especially unscalable. This paper describes what and how IT management processes, technologies and skills can be improved to provide remote customer on-boarding at an appropriate speed for delivery from the cloud. Our contributions consist of major enhancements in a key on-boarding area, namely IT discovery. Experimental results show that our approach aligns the RIM on-boarding methods to the cloud expectations both from a time as well as quality perspective.

1. Introduction

Management services represent a mature subject in the IT business arena. According to Gartner Dataquest, Remote Infrastructure Management (RIM) is a rapidly growing market growing at a CAGR of 36%, and projected to grow from USD $14.3B to $30B by 2010 (Source: Gartner Dataquest, August 2006). Typical RIM service providers use their own Network Operations Centers to remotely monitor and manage customers' IT infrastructure elements such as networks, systems' hardware and operating systems, and applications. The primary business value for RIM services is that it helps global enterprises and SMBs to outsource the burden of managing their IT infrastructure, thus, cutting down costs for infrastructure management and gaining access to expert skills. The customers can focus then on their core business, shifting the responsibility for IT management to RIM, while maintaining ownership of their assets.

A RIM solution generally involves monitoring services comprising of NOC support, reporting, incident notification and escalation, while management services cover problem management and root cause analysis, configuration management, change and release management, maintenance and updates installation. Prior to providing any of these RIM services, the customer has first to select what services to subscribe to during a procedure that is called "on-boarding". Although the IT management itself is rendered more affordable when provided remotely as a service, the RIM customer on-boarding process particularly is not, taking between one to two months of expensive labor.

The current process for RIM customer on-boarding consists of multiple interviews and interactions with customers to 'discover' their IT environment, identify the resources to be managed and guide the enablement of the environment for remote management. This labor intensive approach (measured in weeks) proves to be unscalable when RIM is to be delivered as Management-as-a-Service from an IT infrastructure management cloud.

Cloud computing is an emerging paradigm whereby services and computing resources are delivered to customers over the internet (or intranet) from a service provider who owns and operates the cloud. Cloud-based services characteristically can scale up *promptly* to meet growing demand. The benefit of this will remain unrealized if RIM on-

boarding takes weeks as is the standard today. Since the duration to traditionally provision resources for new RIM customers is comparable to the current on-boarding duration, there is little incentive to motivate change to the current on-boarding approach. However, RIM's goal for delivery from the cloud is 'on-boarding in minutes', which means radical revision of the current approach. To this end, we have identified the following on-boarding problems: (1) lack of a standardized approach or automation for the on-boarding operation flow, (2) inaccuracies in manually assessing the environment from the customer's descriptions or semi-updated inventory files, (3) missing configuration data (e.g., credentials, directory paths, key performance indicators -- KPIs) necessary to setup the monitoring systems, (4) overhead for the SMB customer who is expected to perform complex configurations in their environment (e.g., VPN setup, monitoring data agent/collector installations), (5) evaluated price is not commensurate with the cost of the service expected to be provided.

This paper describes what and how the current IT management processes, technologies and skills can be consolidated to provide an efficient remote customer on-boarding that is appropriate for delivery from the cloud. Our contributions consist of major enhancements in a key on-boarding area, namely IT discovery. The main advantages of using our solution are (i) systematization and increased automation of the on-boarding process, (ii) accurate assessment of customer's environment since its discovery is performed programmatically and confirmed by the customer rather than relying on an inventory provided by the customer, (iii) complete collection of configuration data using automatic look-up and systematic lists of requirements on web site, (iv) acceleration and simplification of the customer experience through reduction of installation and configuration overhead required from the SMB customer. Additionally, the process proposed here also facilitates the setup of a customer single point of contact which accelerates non-technical steps such as the install permissions, provisioning and procurement.

Our experimental results show that the on-boarding approach described in this paper aligns the RIM on-boarding to the cloud expectations both from a time as well as quality perspective.

The remaining of the paper is organized as follows: in Section 2 we detail the traditional RIM on-boarding process and related work in the discovery area; Section 3 introduces the rapid IBM on-boarding process, while the details of our original approach to discovery as a service are described in

Section 4. We show our experimental results in Section 5 and conclude in Section 6.

2. Traditional RIM On-Boarding Process. Related Work

There are many managed services providers in the marketplace. Some are local providers, others regional, and still others global. It is largely the regional and global providers that utilize RIM techniques. They are recruiting IT professionals and making them available to client projects through the use of Global Delivery centers.

For the client on-boarding process, these IT professionals have to identify the client's IT environment either manually during multiple interviews with the customer or by providing a template for exchanging inventory information (e.g., a spreadsheet). Sometimes the customer may provide one by filling out his own questionnaire. However, a better option is programmatic discovery using dedicated discovery software. The inventory information and additional configuration details are then used to configure the monitoring and management toolset. Manual information gathering methods are notoriously error-prone – some of these errors are caught during the tool configuration step which engenders more interactions with the same customer. Other causes that drive the inaccuracy of the manual environment assessment are existing inventory out of date, incomplete or invalid data, and untracked configuration changes (e.g., for the credentials, directory paths, KPIs), that may jeopardize the quality of the RIM service. We will use the manual data gathering performances in the comparison of our experimental results.

When the IT environment discovery is done programmatically, the typical approaches are via stand-alone products, e.g., TADDM [1] or via services that make use of remote product download, e.g., Paglo [2]. Although more accurate in terms of discovery quality compared to the manual approach, the stand-alone products are not suitable for small-and-medium business or SMBs, which have tens to few hundreds of servers. These customers cannot afford nor need sophisticated tools oriented towards large IT enterprises with thousands of IT elements. SMBs prefer to use a streamlined asset discovery service to get the inventory of their IT environment, without the hassle of installing, configuring and managing a discovery product. However, the current remote discovery service providers still require software to be installed and configured by the SMBs in their environment.
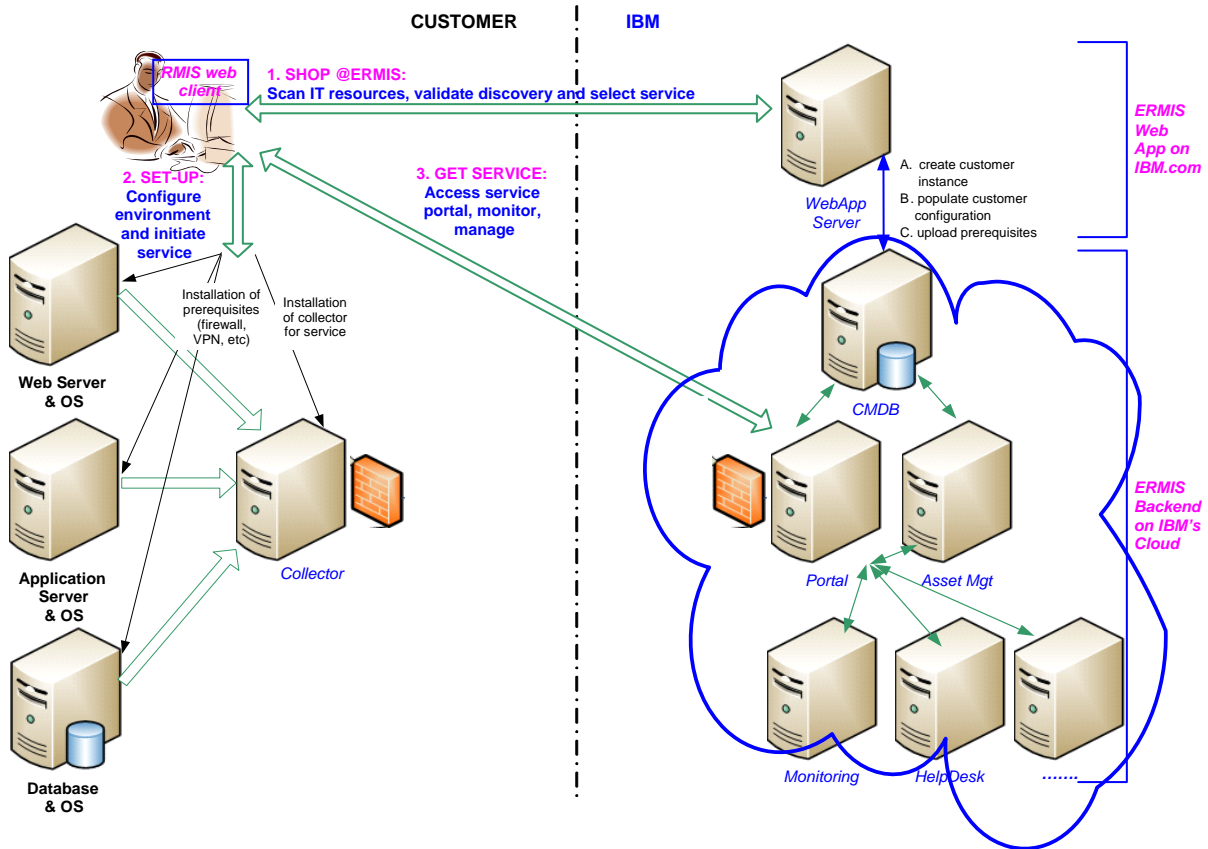
Figure 1. Enhanced on-boarding process for RIM delivered from an IT management cloud.

In [2] the provider offers the discovery tool for free in the context of their monitoring service and the customer has to take care of the discovery tool installation. This is due to the fact that the discovery must take place from a node in the network to be probed, to circumvent firewalls, Network Address Translation (NAT), and other impediments. Many users, especially in the SMB space, prefer not to face the burden of this administrative overhead and may lack the necessary skills required. After the discovery is completed, the IT professionals have typically to manually collect additional configuration details since the discovery provides partial detection of the environment. We call this approach **manual on-boarding with automated discovery** and will compare in Section 5 its characteristics to our approach presented in this paper as well as to the manual approach.

Other related solutions that involve remote discovery include [3] which uses a browser to control a discovery process, however, it is unclear whether the provider intends for the "NDM Agent" to be running on the web server or some other machine (or whether it should or could be located in the browser). They also provide "passive discovery" which

involves packet sniffing to discover applications running on client machines. JLocator [4] describes a Java applet based network discovery tool. They discover the network topology, but do not examine applications neither services on the identified elements. XAssets [5] is a service comparable to SNAPPiMON [6], **Error! Reference source not found.**where the management tool is also browser based. In [5], the discovery process uses a wide variety of discovery techniques to collect hardware items details, while unrecognized software items lists from customers are sent on a regular basis to the provider staff and these items are manually investigated and added to the discovery database. In [6] the discovery is also a combination of manual discovery, for network and server level items and credentials, and automatic look-up for OS and application configuration.

Once the inventory has been discovered and validated, and the additional information on the resources to be managed gathered, the IT professionals proceed to or guide the customer through the enablement of the environment for remote management. This step consists of the installation of data collector or agents into the
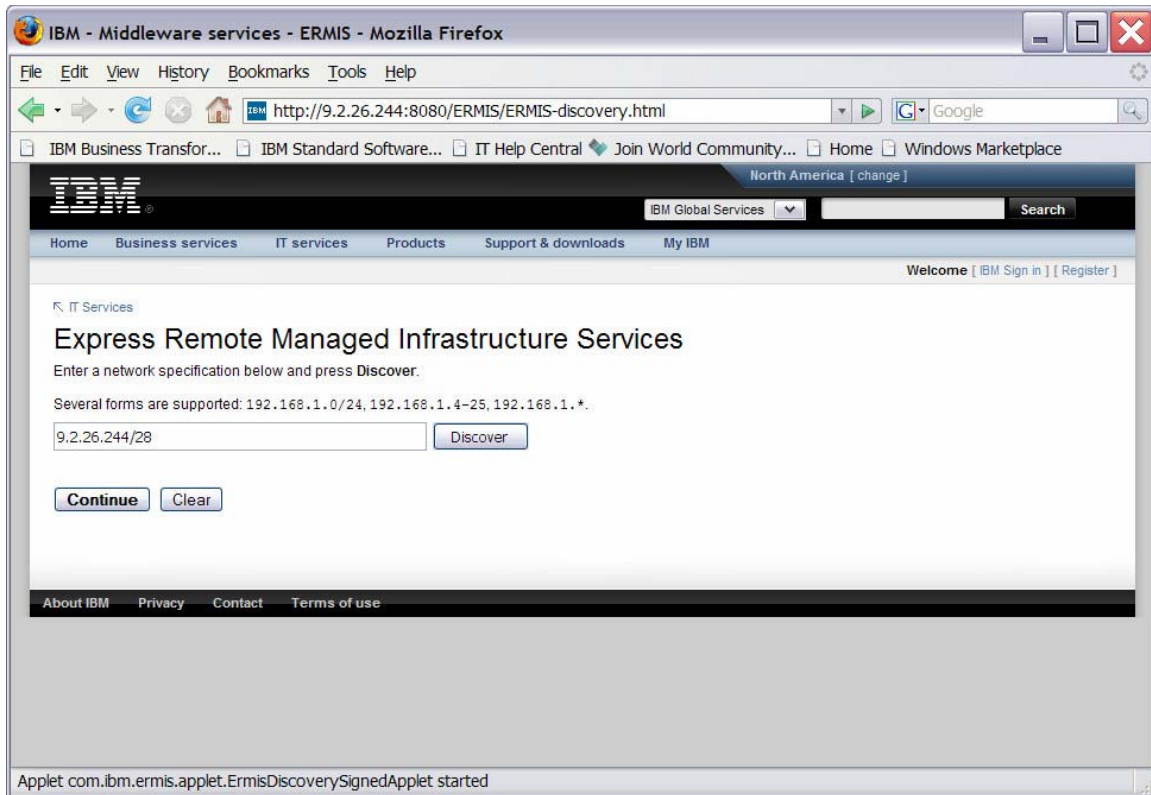
Figure 2. Applet prompts customer for network address range to scan.

customer's premises, firewall configuration for site-to-site VPN set-up and NATing of endpoints. Finally, upon performing all necessary data collection and setup for monitoring and managing the selected items in the customer's environment, the RIM provider prices the offering and starts delivering the IT management service.

In the following section we will present our enhanced process for cloud delivered RIM on-boarding and detail the discovery aspects of the process. The collector/agents installation, firewall configuration for site-to-site VPN set-up and NAT-ing of endpoints will make the subject of a different publication.

3. IBM's Rapid On-Boarding from Cloud

Efficiency is central to delivery from a cloud since most of the cloud benefits derive from the speed and simplicity with which services can be put into production and reprovisioned. We describe a highly automated RIM on-boarding process where the main objective is to circumvent the limitations of the on-boarding process detailed above. The main steps are as follows:

1. SHOP: The customer shops for remote management on a RIM specific web application at the service provider's sales portal (in our example this is www.ibm.com), using a browser on local machine. The RIM web application downloads and launches via a RIM web client an IT *discovery* tool to detect the customer's environment. The IT elements discovered are presented on the browser as soon they are identified. The SMB customer can update and thus *validate* the findings in accordance with their environment. From the resources validated on the web application, the SMB customer can then *select* those that are to be monitored and managed in their environment. The RIM web application sends the inventory and configuration data back to the RIM back-end on the cloud, where a new customer instance is created on the Configuration Management Database (CMDB) and the related inventory and assets are populated with the information collected from the customer. The RIM back-end computes the service price and prompts the customer to "Buy" via the browser.

2. SET-UP: Upon selection of "Buy", the RIM web application prompts the customer to *provide* additional information (credentials, configuration, etc.) necessary for enabling monitoring on the selected machines. The RIM customer proceeds with
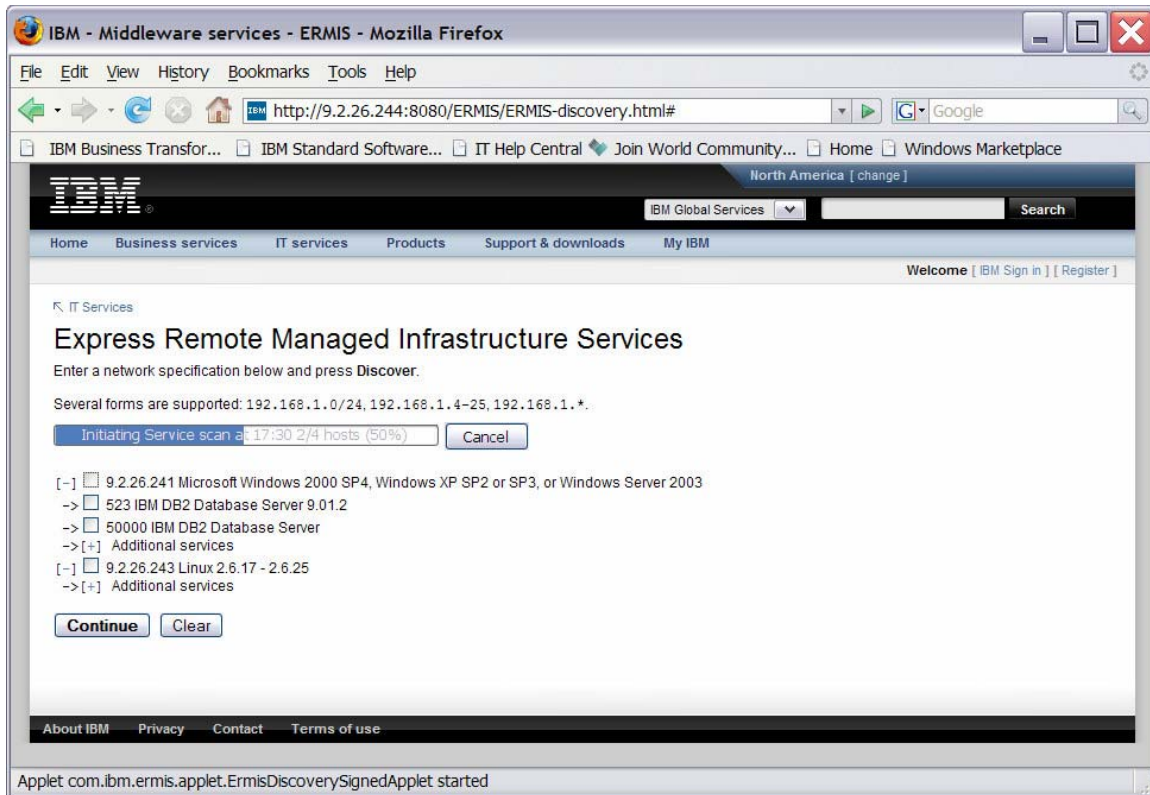
Figure 3. As hosts and services are discovered, they are presented in the web page.

filling out forms requiring additional configuration items for those systems selected for monitoring. Either the RIM web application or the back-end on the cloud checks for all the required monitoring pre-requisites for all the selected systems, middleware and applications, and tries to seamlessly install and configure the prerequisites. The prerequisites that cannot be taken care of automatically are to be manually installed.

Then the customer initiates the firewall configuration for site-to-site VPN and NAT-ing of endpoints with minimal or no customer intervention. The RIM customer also *installs* and *configures* the necessary software (agents, collector, etc., depending on the RIM techniques employed) to enable monitoring and management at the customer premises. The RIM web application provides as much as possible automation and guidance on this step.

3. GET SERVICE: The RIM back-end in the cloud begins the remote monitoring and management of the customer's IT environment. The RIM back-end sends invoice and collects payment.

The automation leveraged in our on-boarding process described here takes advantage of web techniques that allow the reliable replacing of a significant number of manual steps in the traditional process. We detail next these techniques for the discovery step of the on-boarding, which is the focus of our contributions in this paper.

4. Discovery

IT discovery service requires direct access to the customer's network. This level of access is generally not available to web applications, which run mainly in a hosted environment on a remote server in a data center. The web browser mediates all interaction with the customer's host system and prevents exactly this type of access for security reasons. Existing discovery tools require installation, configuration and use of software onto the user's machine. These steps break the discovery service automation and incur an burden on a potential customer who may loose interest in such a RIM offering.

Our system provides an automated *credential-less* IT asset and application discovery mechanism in the web application which allows the service provider to begin the RIM service quickly and accurately. Without such a tool, the process of extracting accurate information about the customer's systems is tedious and error prone. By integrating this service into the shopping process and web application, the
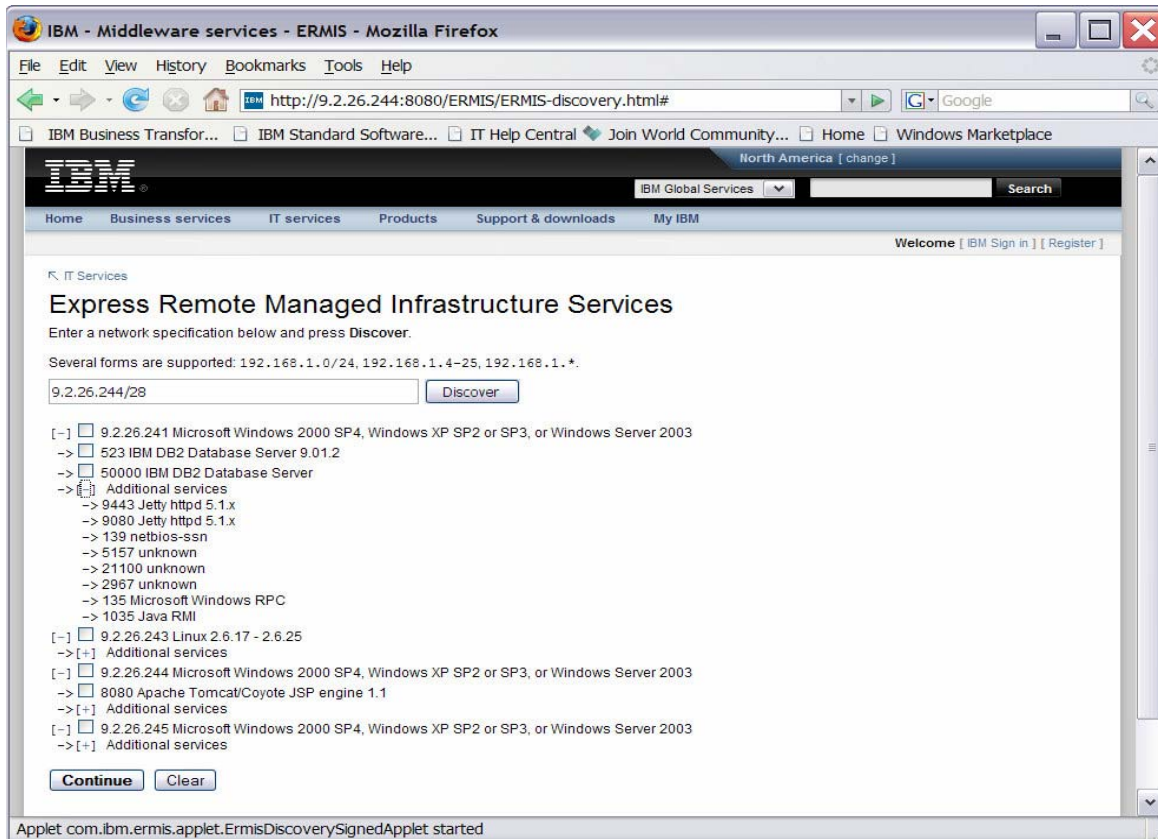
IBM - Middleware services - ERMIS - Mozilla Firefox

File   Edit   View   History   Bookmarks   Tools   Help

http://9.2.26.244:8080/ERMIS/ERMIS-discovery.html#

Google

IBM Business Transfor...   IBM Standard Software...   IT Help Central   Join World Community...   Home   Windows Marketplace

North America [change]

IBM Global Services   Search

Home   Business services   IT services   Products   Support & downloads   My IBM

Welcome [IBM Sign in] [Register]

IT Services

Express Remote Managed Infrastructure Services

Enter a network specification below and press **Discover**.

Several forms are supported: 192.168.1.0/24, 192.168.1.4-25, 192.168.1.*.

9.2.26.244/28   Discover

[−] □ 9.2.26.241 Microsoft Windows 2000 SP4, Windows XP SP2 or SP3, or Windows Server 2003
 −> □ 523 IBM DB2 Database Server 9.01.2
 −> □ 50000 IBM DB2 Database Server
 −> [−] Additional services
   −> 9443 Jetty httpd 5.1.x
   −> 9080 Jetty httpd 5.1.x
   −> 139 netbios-ssn
   −> 5157 unknown
   −> 21100 unknown
   −> 2967 unknown
   −> 135 Microsoft Windows RPC
   −> 1035 Java RMI
[−] □ 9.2.26.243 Linux 2.6.17 - 2.6.25
 −>[+] Additional services
[−] □ 9.2.26.244 Microsoft Windows 2000 SP4, Windows XP SP2 or SP3, or Windows Server 2003
 −> □ 8080 Apache Tomcat/Coyote JSP engine 1.1
 −>[+] Additional services
[−] □ 9.2.26.245 Microsoft Windows 2000 SP4, Windows XP SP2 or SP3, or Windows Server 2003
 −>[+] Additional services

Continue   Clear

Applet com.ibm.ermis.applet.ErmisDiscoverySignedApplet started

Figure 4. Automatically discovered services and hosts can be selected for monitoring and management

customer is able to know exactly which systems and applications are on the customer's internal network, select them for remote service, and receive an immediate offering price. Credential-less discovery performs its task without requiring authentication credentials to access remote systems. Hence, the discovery must be done in a *black box* manner, e.g., on an operational basis by examining how the remote systems respond to probes sent over the network. The probes interact only with services that are actively listening on the systems' ports. This means that the discovery tool can only detect applications which provide services via the network and which support a protocol for which a probe has been written. Most desktop applications cannot be discovered in this way; however, a wide range of applications and services can be detected and interrogated. Indeed, applications that are of interest to customers of RIM services are of this type: web servers, application servers, database management systems, mail systems, and many other classes of applications can be detected reliably this way.

Our solution provides a signed Java applet which downloads and runs an existing open source tool called Network Mapper (Nmap) [7] to perform the discovery. By signing the Java applet with a key that has been certified by a trusted authority, the applet can escalate its privileges to execute Nmap which then can access the network to perform the port scan, OS detection, and service discovery. We further maintain the appearance of discovery-as-a-service by using the applet just to collect the discovery results. As it retrieves host and service information from Nmap, it posts this information to a Java servlet on the web server. This data is then conveyed to the user's browser via AJAX-style requests using JSON-encoded data objects. The information is then presented inside the web page using Javascript, maintaining the appearance that all the data has been collected and is being provided by the web service.

Figures 2 through 4 show the discovery process through a series of screen captures. When the applet is initially started, it downloads and extracts the low level scanning tool. Once this is complete, it provides a prompt which allows the user to enter a range of IP addresses (Figure 2), defaulting to the networks to which the host is currently attached. Once discovery

Table 1. Comparison of on-boarding methods using manual inventory collection, automated discovery and discovery-as-a-service

	Manual on-boarding inventory collection	Manual on-boarding with automated discovery	Discovery-as-a-Service
Time to gather inventory data	10 days	2-3 days	1-3 hours
Accuracy of results	Poor	Good	High
Labor cost	High	Medium	Low
Software cost	Low	High	Low
Service Quality	Satisfactory	Excellent	Outstanding

has commenced, hosts and services are presented to the user as they become available from the low level discovery tool (Figure 3). The user can interactively browse the results at any time. When complete, the user can select the hosts and applications to monitor and continue (Figure 4). Once the end points have been selected, they are stored into the CMDB.

5. Results

The results presented in this section compare the three discovery approaches described above, i.e., fully manual and manual using a stand alone product (see Section 2) and discovery-as-a-service (see Section 4). Recall that the fully manual inventory collection process represents the typical on-boarding process where the customers enter their systems and applications into a spreadsheet which is passed between the customer and the service provider as errors are found and updates are required. Discovery using a standalone product, such as TADDM or Nmap, replaces the manual step of initially entering the end point information such as IP address, operating system, and application names and ports with the step of running Nmap in the customer's environment – typically done by an associate of the service provider that visits the customer site, bringing a discovery appliance which can scan the network for this information. In the discovery-as-a-service process when the customer shops for RIM services at a web site, an web application integrated discovery

tool provides the remote discovery, allowing the user to extract reliable host information within an hour or so (the time depends on the size of the network). The user selects the hosts to be monitored, providing additional information and credentials necessary for monitoring and management, and proceeding to enter billing information and begins using the service.

Since it is difficult to collect hard numbers to compare these different processes, we present here qualitative average results based on our experience with RIM on-boarding processes and tools. Table 1 illustrates the advantages of on-boarding a new RIM customer using a discovery-as-a-service platform. Manual on-boarding inventory collection, even when a traditional discovery tool is used, can take several days, whereas with our web based approach it can be completed in a couple of hours.

Because the fully manual process requires negotiation even about which machines are present, the accuracy of the results is poor; the customer frequently will not know all the systems and the services running in his environment. The use of a traditional discovery tool by an IT professional during the on-boarding process improves the quality of the scan because it will find the endpoints that are responding at the time of the scan. Still, with this approach, the customer may decide not to select all endpoints to be monitored. Similarly, there may be additional services to be monitored, which were not present or discoverable. Hence, there is again a need to negotiate. When discovery is provided as a service,

the customer can select endpoints for remote service as they are discovered, as well as correct or append any information early in the process. This leads to significantly high accuracy of data collected for service delivery, to reduced negotiations and thus to an improved service quality.

From the cost perspective, due to the fact that manual on-boarding is time intensive, requiring multiple rounds of negotiation about which systems to monitor and how to access them, the labor cost of providing initial support to the customer generates a high on-boarding cost. Similarly, when an IT professional is deployed to a customer environment with discovery software, the cost may be even higher, due to software licensing costs as well as the IT professional labor dedicated to a single customer. When a RIM service is provided from the cloud together with a discovery service, these per-customer costs are reduced dramatically, since the software costs are shared and the majority of customers do not require on-site support at all.

6. Conclusions and Future Work

We have presented a novel solution to the problem of rapid on-boarding for customers of remote managed infrastructure services. The approach involves providing credential-less asset discovery as a service on the web to catalog the hosts, operating systems and service applications. We have detailed solutions to the challenges that arise from providing the discovery service: 1) using a minimal signed applet to perform actions which require elevated privileges on the computer from which the customer

is shopping, and 2) improving the cost, quality, time, and interactivity of the scan process.

We continue to study existing remote managed services and prototyping capabilities that would enable them to be offered as a cloud-based offering. This includes looking at automating site to site VPN connectivity, alternate protocols to do touchless remote monitoring of customer environment from the cloud, automated service provisioning and a Service Request Catalog to make it a standardized offering with fixed and structured service request choices from a services catalog on a metered basis.

References

[1] Tivoli Application Dependency Discovery Manager (TADDM)   - http://www.ibm.com/software/tivoli/products/taddm

[2] Paglo, http://www.paglo.com/opensource/paglocrawler

[3] Design of Hybrid Network Discovery Module for Detecting Client Applications and ActiveX Controls, http://www.springerlink.com/content/y51p5g76k25578g1/

[4] JLocator: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.39.5030

[5] XAssets: http://www.xassets.com/discovery.aspx

[6] SNAPPiMON   -   http://www.snappimon.com/, Copyright IBM Corp.

[7] Nmap - http://nmap.org/