



# ESG Lab Validation Report™



## EMC Smarts Application Discovery Manager Automated Infrastructure Discovery

A validation study  
by  
ESG Lab  
July 2007

Authors:  
Mark Bowker, Brian Garrett, Bob Laliberte

## Table of Contents

<b>Introduction .....</b>	<b>3</b>
<b>ESG Lab Validation .....</b>	<b>4</b>
<i>Getting Started.....</i>	<i>4</i>
<i>Application Discovery .....</i>	<i>5</i>
<i>High Definition Discovery.....</i>	<i>7</i>
<i>Application Dependency Mapping.....</i>	<i>10</i>
<i>Reporting .....</i>	<i>12</i>
<b>ESG Lab Validation Highlights .....</b>	<b>13</b>
<b>Issues to Consider.....</b>	<b>13</b>
<b>ESG Lab's View .....</b>	<b>14</b>

### ESG Validation Reports

The goal of ESG Lab Validation reports is to educate customers about various storage and storage-related products, including storage systems, backup-to-disk solutions, storage management applications, backup and recovery software, storage virtualization platforms, etc. ESG Lab reports are not meant to replace the necessary evaluation process that end-users should conduct before making purchasing decisions, but rather to provide insight into these technologies. Our objective is to go over some of the highlighted features/functions of such products, show how they can be used to solve real customer problems, and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This report was commissioned by EMC Corporation.

## Introduction

EMC acquired privately held nLayers, Inc. in July 2006 for the foundational technology that is currently known as Smarts Application Discovery Manager (ADM). Smarts ADM continuously discovers and monitors application components, resource dependencies, service levels and usage models within data centers. This ESG Lab validation report examines the core capabilities of Smarts ADM and version 5.0 features, which became generally available in October 2006 and includes a new dashboard design and high definition discovery.

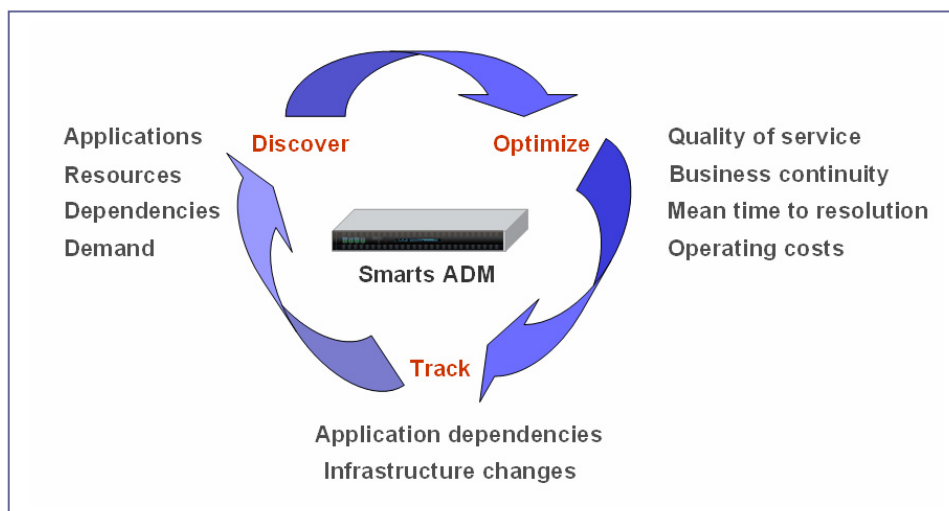
## Background

Managing a data center at the beginning of the 21<sup>st</sup> century is a daunting task. Complexity is inherent in large data centers. The process of identifying all of the applications and the infrastructure they rely on can be a full time job. Applications are often interdependent, which exacerbates the complexity. Trying to identify and track these interdependencies in an ever changing environment is almost impossible to do manually. Once all the information has been collected and correlated, change often renders the information out of date and useless.

Now, more than ever before, managers rely on IT services to run their business and fuel growth. Applications need to be managed to very high levels of service and interruptions are often unacceptable. Yet change is inevitable. Mergers and acquisitions, data center consolidations, server consolidations and maintenance activities need to be managed, while mission critical applications remain on-line and available. To better manage the inevitability of change, IT managers need tools capable of tracking the interrelationships between applications, servers and infrastructure.

About two years ago, EMC embarked on a path aimed at providing its customers with intelligent, actionable information spanning the data center infrastructure space with their purchase of Smarts. EMC has steadily increased the scope of the Smarts product line—known for its root cause analysis and incident management code book technology—by expanding into storage and applications. Application Discovery Manager (ADM) plays an important role in the Smarts product line. Smarts ADM provides automated discovery of the dependencies between applications and resources, which can be used to optimize and track dependencies with a goal of extracting real business value from the IT infrastructure.

Figure One: Smarts Application Discovery Manager



This ESG Lab report examines Smarts ADM version 5.0 capabilities, including agent-less passive discovery, automated fingerprinting of applications, an open reconciliation engine, a new and improved Discovery dashboard, authenticated high definition discovery and topology views and reports, which map the dependencies between applications and servers.

## ESG Lab Validation

ESG Lab performed a hands-on evaluation of EMC Smarts Application Discovery Manager at the EMC laboratories located in Hopkinton, Massachusetts and Santa Clara, California. The evaluation began with a high level overview of the automated agent-less architecture used to discover applications deployed at each location.

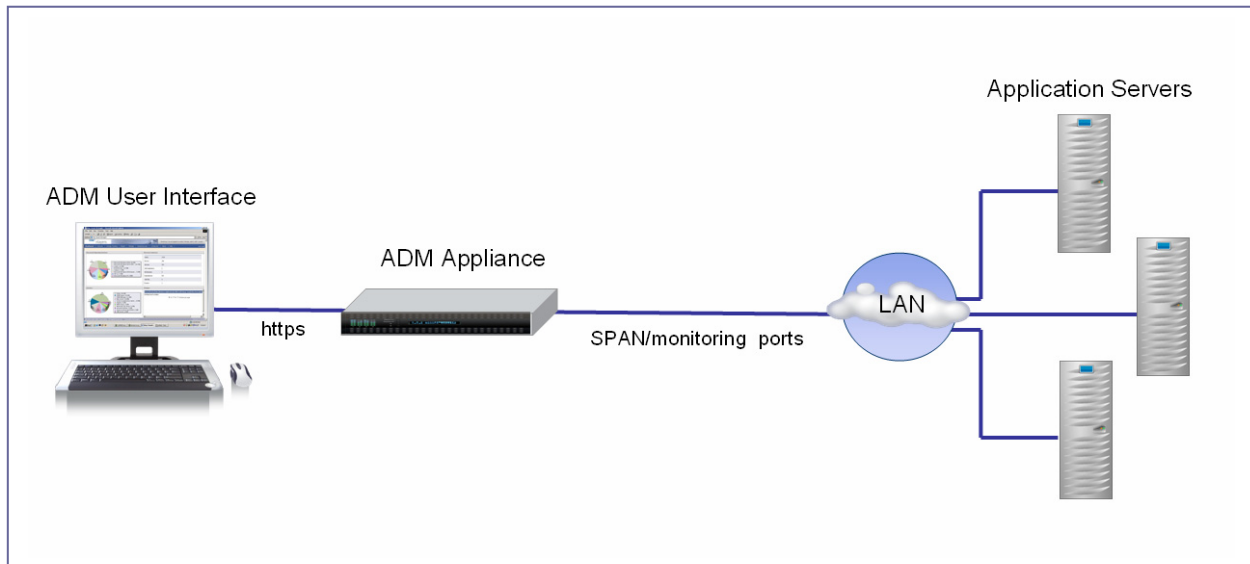
### Getting Started

Smarts ADM is shipped as a self-contained, rack-mountable appliance with software pre-installed. As shown in Figure Two, the Smarts ADM appliance monitors application networking traffic on a corporate network using switch monitor ports. IT managers access the Smarts ADM user interface from a web-based console over a secure https connection.

Smarts ADM uses monitoring ports to passively observe the traffic that goes through Ethernet switches without modifying packets or sending new packets through the network. A variety of industry standard switch traffic monitoring protocols including SPAN, VSPAN and RSPAN are supported. Using deep packet inspection on a sampled basis, Smarts ADM monitors applications running over the corporate network. There is no need to install agent software on servers.

A Smarts ADM appliance can be connected to one or more aggregation switches. Each aggregation switch provides the ability to connect to 32 switch monitor ports. A single Smarts ADM appliance can be deployed in less than one hour to monitor the network traffic between thousands of servers and applications.

**Figure Two: Deploying Smarts ADM within a Corporate Local Area Network**



### Why This Matters

IT managers are struggling to keep up with the number of servers and applications being deployed and the dependencies between them. Automated monitoring to better understand what's running where would clearly save time and money, but at what cost? Traditional network-based methods that scan or probe the network add network overhead and provide little application-level insight. Host-based methods that provide insight at the application level typically require the use of agents. Host-based agents are hard to manage, consume server resources and introduce risk. ESG Lab has confirmed that Smarts ADM provides agent-less application level insight with zero network overhead.

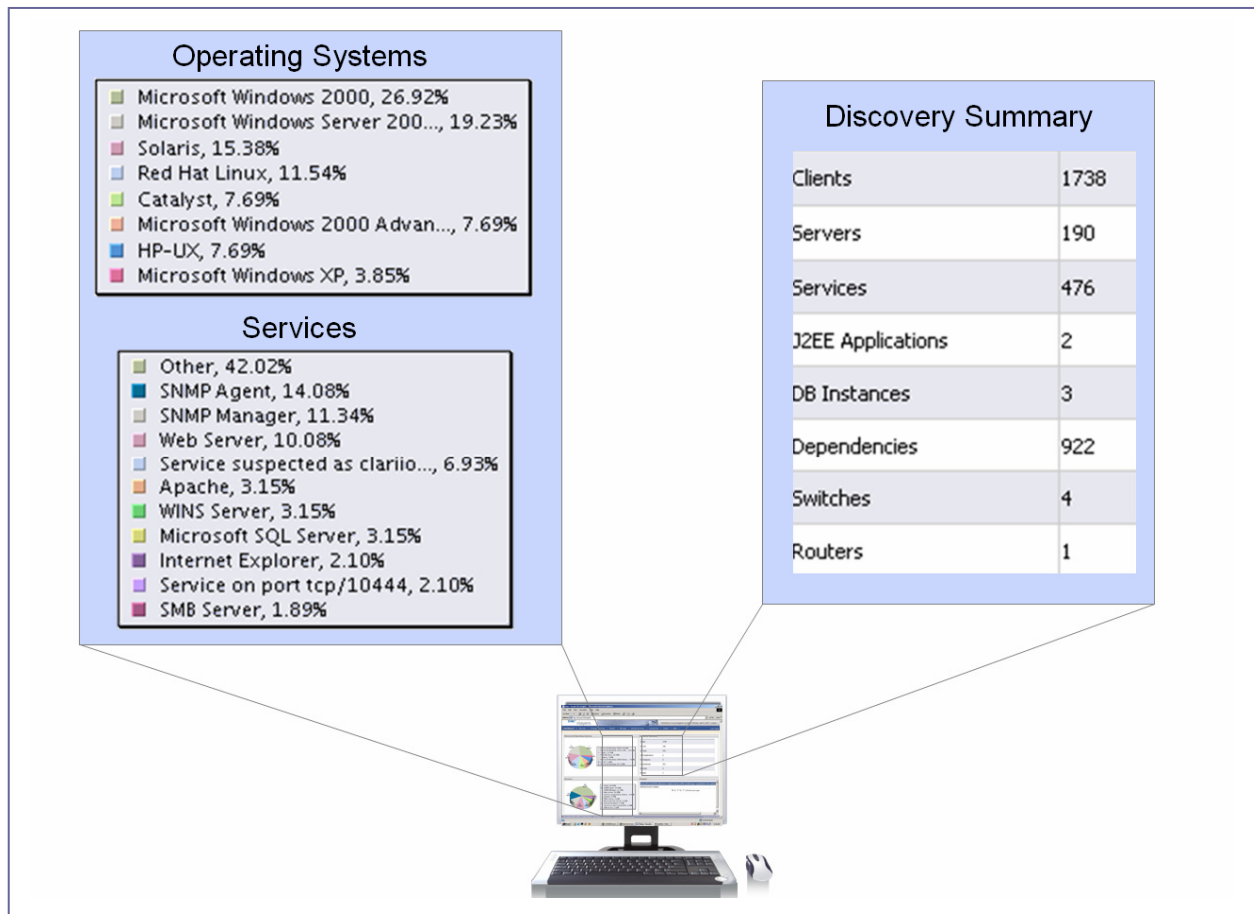
### Application Discovery

A Smarts ADM appliance uses deep packet inspection on a sampled basis to discover all of the servers and applications deployed within a corporate network. Heuristic fingerprinting algorithms are used to recognize application interactions and map the behavior against a list of known applications. Redundant information collected by the appliance is eliminated by an open reconciliation engine before being committed to a common management repository that can be fed into a customer’s chosen CMDB. For targeted audit and compliance activities, many customers use ADM’s management repository on a standalone basis, too.

### ESG Lab Testing

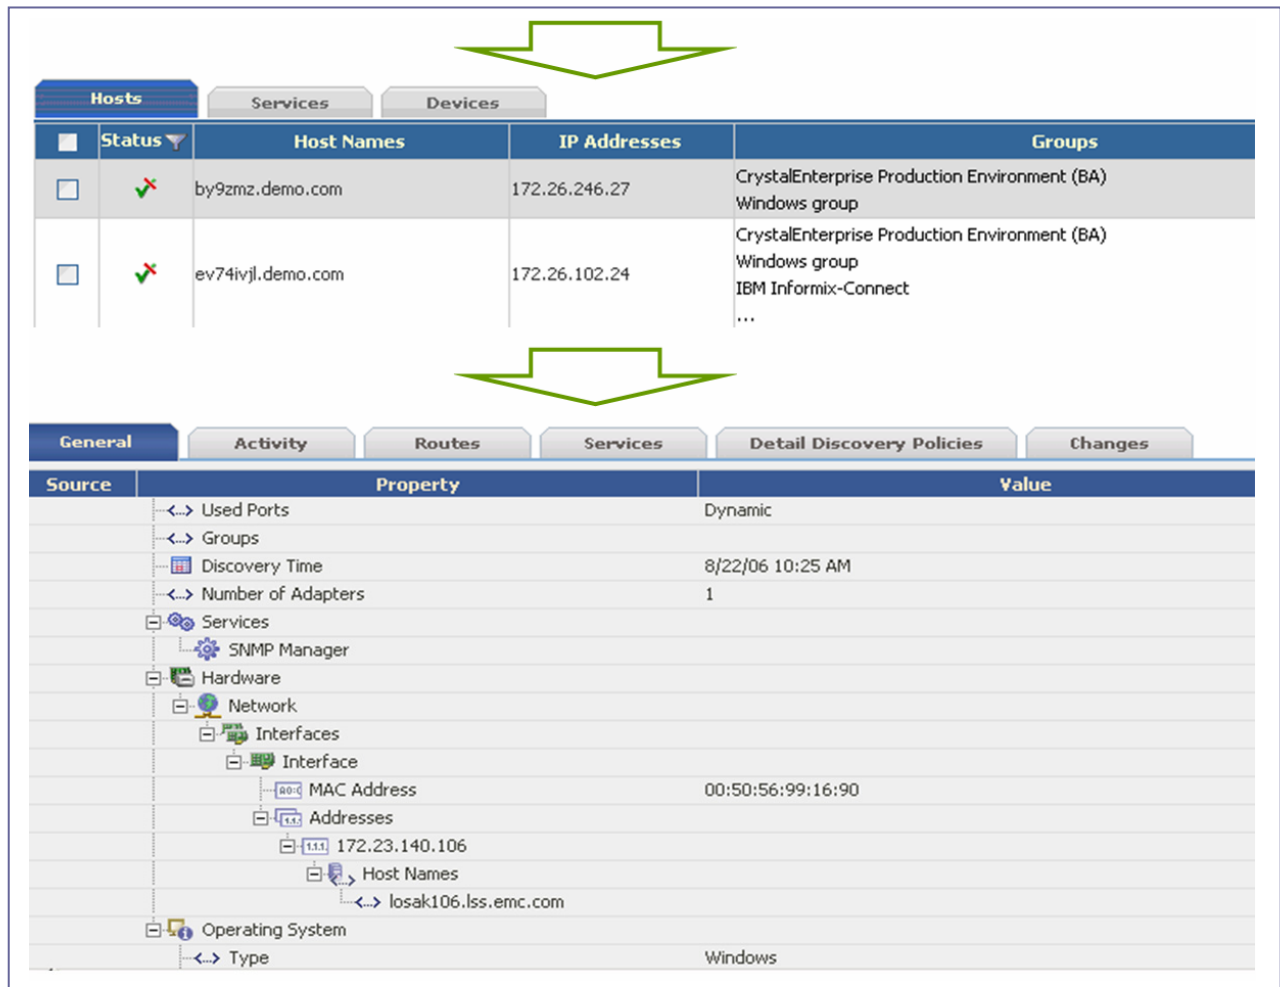
ESG launched the web-based management console to view application discovery results. As shown in Figure Three, the home page presents a comprehensive dashboard level summary of the environment. Note the number and variety of operating systems, services, clients and servers discovered automatically by Smarts ADM.

Figure Three: The Smarts ADM Dashboard



The Smarts ADM management console not only provides a high level summary of the application environment, but also provides a variety of tools which can be used to drill down and learn more about servers, applications and the interdependencies between application services. ESG Lab began a deeper investigation of the application infrastructure using a host inventory view as shown at the top of Figure Four. Drilling down further using the General tab for a particular server, ESG Lab was impressed by the level of detail collected by Smarts ADM including the operating system and applications detected on each server.

Figure Four: Drill-down Application Discovery



### Why This Matters

An inventory of the servers, operating systems and application-level services deployed within a corporate network provides extremely valuable insight, which can be used to maximize resource utilization, reduce infrastructure and licensing costs and accelerate problem resolution. With traditional methods, this is a difficult task because you first need to know what you have before you can perform an inventory. ADM automatically finds all of the servers, applications and switches in your environment—even those that you didn't know were deployed or were long forgotten.

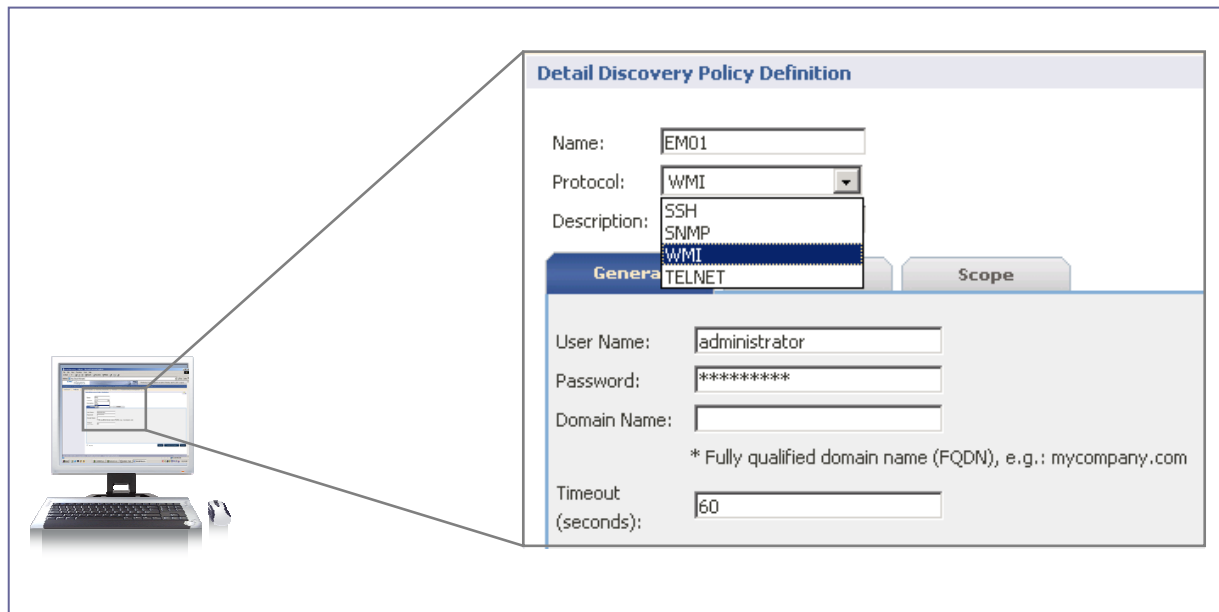
### High Definition Discovery

High Definition Discovery (HDD), a feature introduced in 2006 in Smarts ADM version 5.0, provides the optional ability to gather more in-depth information than the sampled deep packet inspection method described earlier. This deeper level of discovery is performed using industry standard maintenance interfaces like the Windows Management Interface (WMI), secure shell (ssh), the simple network management protocol (SNMP) and telnet. HDD is used to gather more detailed information about the underlying infrastructure supporting applications and the dependencies between application components.

### ESG Lab Testing

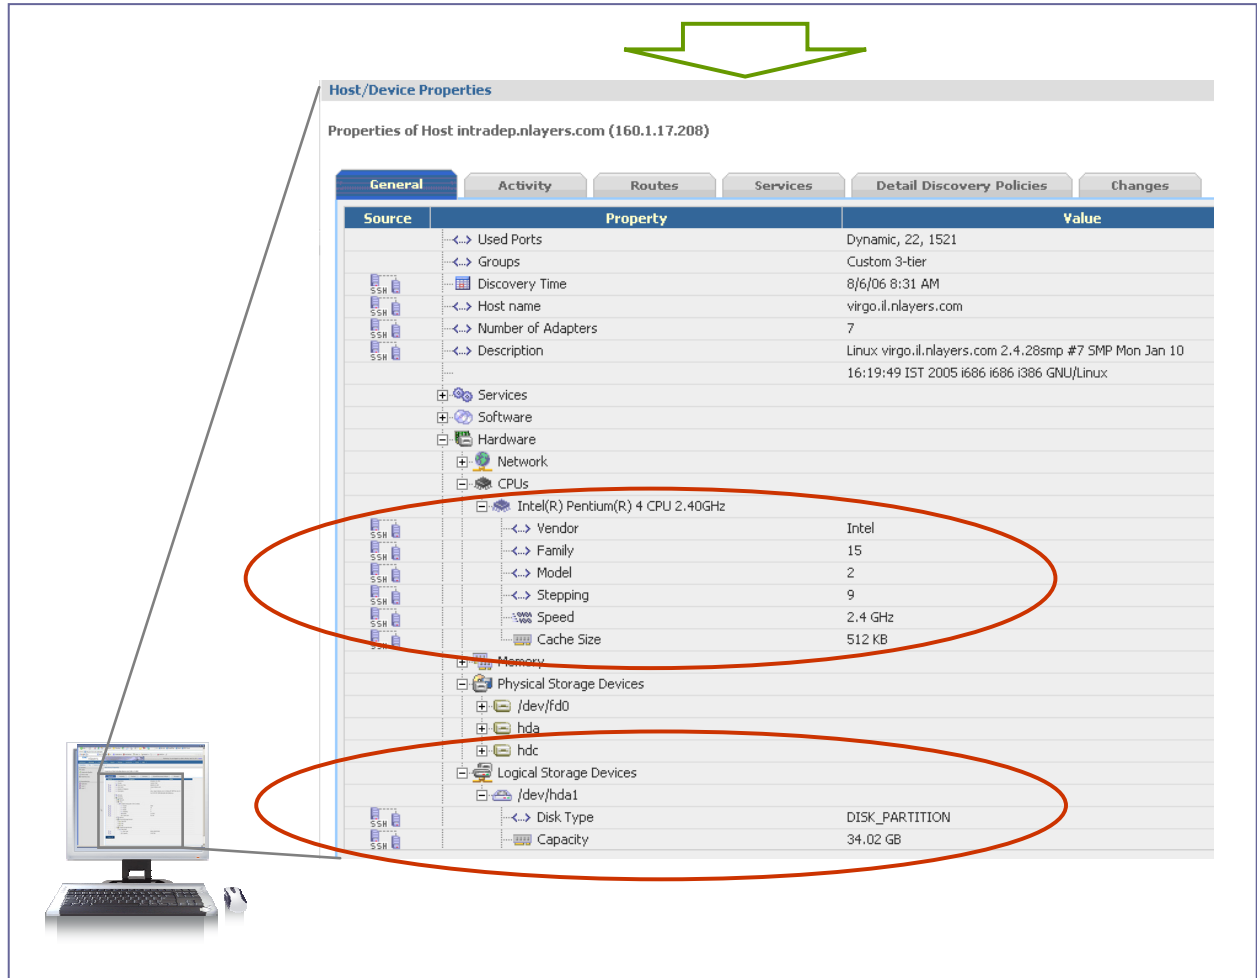
Access to maintenance interfaces like WMI is typically restricted to authorized administrators. As such, before high definition discovery data can be collected, Smarts ADM must first be configured with a valid set of authorized user credentials. ESG Lab configured a Windows server for high definition discovery as shown in Figure Five. The process was straightforward – simply select the network name of the resource to be managed (in this case server EM01), the protocol to be used for high definition discovery (in this case WMI) and a user name and password.

Figure Five: Configuring Credentials for High Definition Discovery



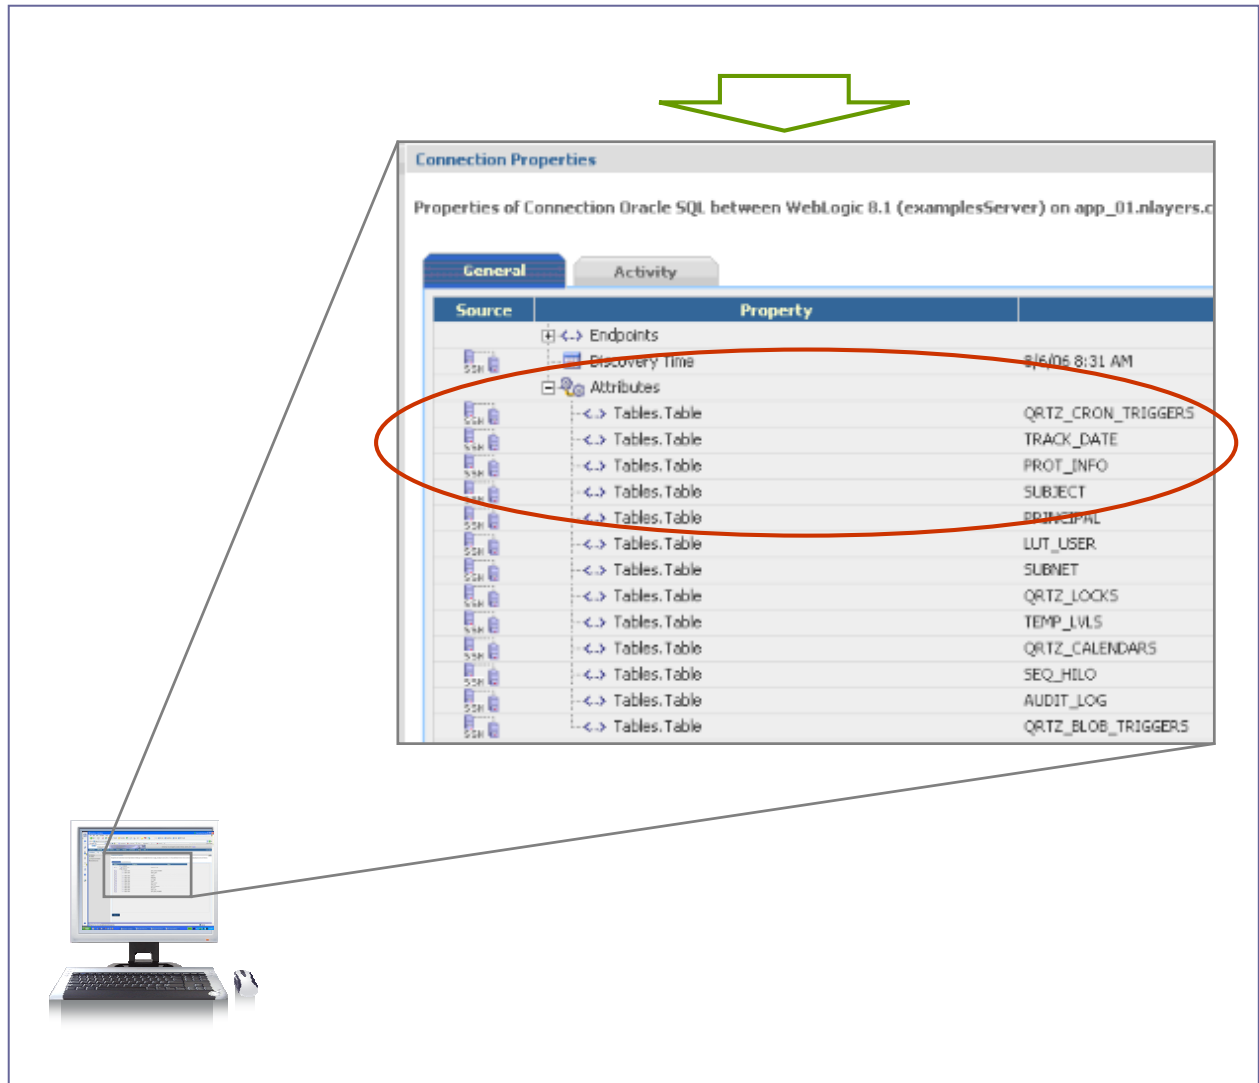
ESG Lab used High Definition Discovery to drill down further and learn more about the application infrastructure in the lab at EMC. Evidence of this new level of detail was first noticed as Smarts ADM displayed low level server, operating system and application level details as shown in Figure Six. This screen shot was taken as ESG Lab investigated processor and logical storage details. Note how an icon towards the left is used to denote that the source of the high definition discovery information. In this case an ssh icon is shown for a Linux server.

Figure Six: High Definition Discovery



Drilling down even further, ESG Lab was impressed by the level of detail captured and distilled during High Definition Discovery. Java 2 Enterprise Edition (J2EE) models and components and their relationship within and between components on other servers was observed. Detailed information at the application level for common applications including Oracle and Microsoft Exchange was noted. For example, the screen shot shown in Figure Seven shows the table level details collected and presented by High Definition Discovery while Smarts ADM observed SQL traffic between a BEA WebLogic application and an Oracle database server.

Figure Seven: Oracle SQL Server High Definition Discovery



### Why This Matters

IT managers need a solid understanding of the applications environment and how configuration changes and problems in one domain might impact the multiple layers of the IT infrastructure. Without this knowledge, introducing planned changes is a gamble. Gathering this information manually is resource intensive and error prone. ESG Lab has confirmed that real-time High Definition Discovery can be used to eliminate a reliance on guesswork while minimizing change-induced service outages, degradations and failures.

### Application Dependency Mapping

Keeping track of the dependencies between services, applications and servers can be a frustratingly difficult task for IT managers. This information is needed to answer a variety of seemingly simple questions including:

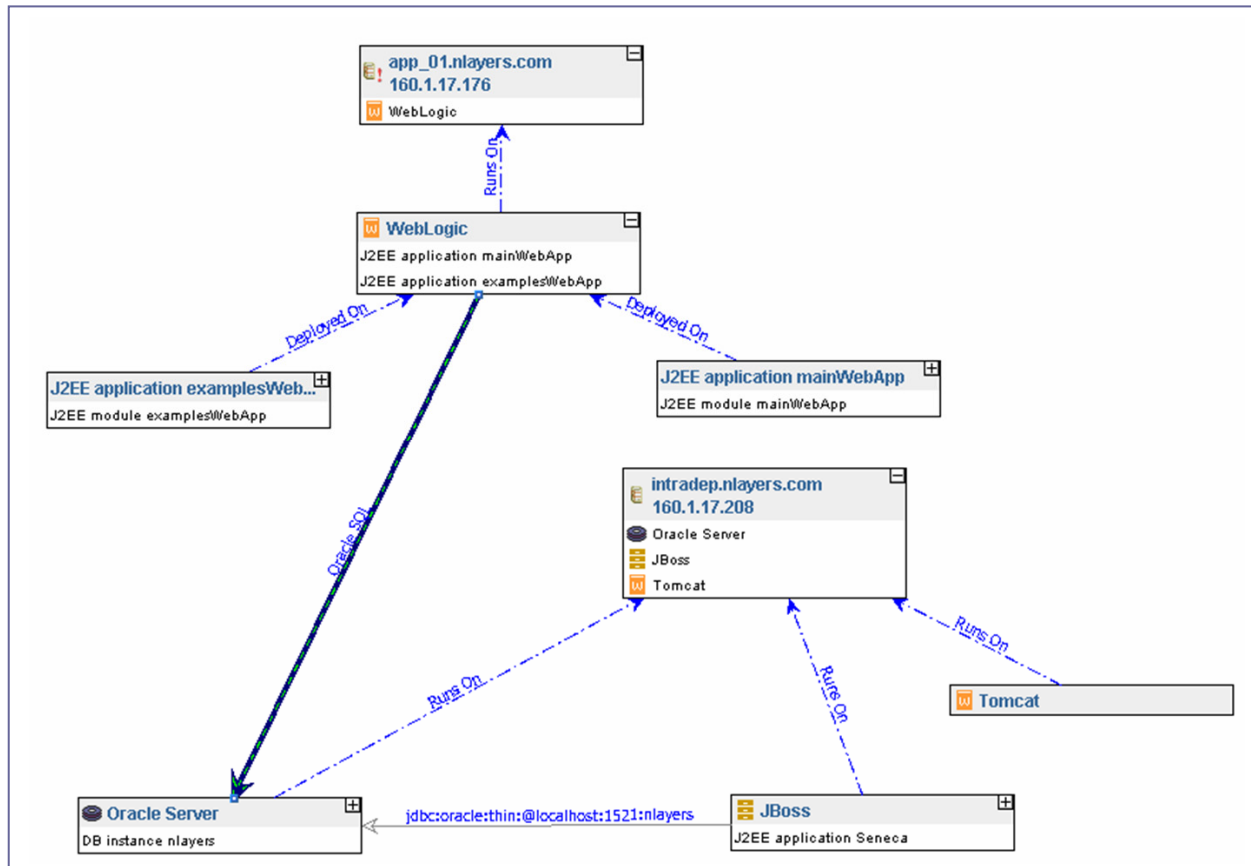
- Which applications and users will be impacted if I take down this server for a software upgrade?
- What needs to be replicated to keep this mission critical application running during a disaster?
- If I take down this server for a hardware upgrade, which applications will be effected?
- When we move to the new data center, which applications and services need to migrate together?

Smarts ADM not only tracks the applications running on each server, but also keeps track of the dependencies between application services running on one or more servers. Smarts ADM provides a topological view of those dependencies by monitoring the network traffic between services and the knowledge gained during deep packet inspection.

### ESG Lab Testing

ESG Lab began the investigation of application service dependencies by launching the Discovery Map from the management console. As shown earlier in Figure Three, Smarts ADM detected 922 dependencies between application services running on 190 servers and 1,738 clients in the EMC lab in Hopkinton. When first viewed at a high level, the topology map depicting all of the 922 dependencies was dauntingly complex. Zoom and isolation features in the topology viewer were used to drill down to view a manageable subset of the topology. With guidance from EMC, a more practical means of navigating the topology was pursued by working from the bottom up, starting with an Oracle application running on a particular server instead of working from the top down starting with the entire data center.

Figure Eight: Smarts ADM Discovery Map



Starting the exploration from the Oracle application server, a drill down from the host inventory viewer was used to discover its dependent services. A series of drill downs and right clicks were used to visually explore the services and dependencies associated with the Oracle server, including a BEA WebLogic application running on another server.

### How to Read the Map

The discovery map shown in Figure Eight depicts three powerful and useful levels of dependencies which were detected by Smarts ADM software after sampling the network traffic between application services:

1. **Application to services relationships**  
The Oracle Server, Java J2EE and Apache Tomcat services associated with the Oracle Server application are running on the server named intradep (160.1.17.208).
2. **Server to server relationships**  
The WebLogic application services running on the server named app\_01 (160.1.17.176) rely on the Oracle application services running on the server named intradep (160.1.17.208).
3. **Service to service relationships**  
The WebLogic Java J2EE service is communicating with the Oracle Server using the Oracle SQL protocol.

### Why This Matters

The only constant in business and Information Technology is change. Understanding the relationships between infrastructure and application services is vital as IT managers react to a constant barrage of changes including mergers and acquisitions, data center consolidations, server consolidations, application upgrades, server maintenance, lease renewals and bugs. Figuring out which services are associated with an application is hard. Figuring out the dependencies between applications and services spread over many servers is even harder. This information is unique for each organization and is often maintained in out of date spreadsheets, documents and diagrams. ESG Lab has verified that ADM provides valuable real-time maps which can be used to learn the dependencies between application servers while accelerating the ability to react to change with confidence.

## Reporting

Smarts ADM includes a rich set of reporting tools that can be used to provide real-time reporting on application infrastructure. Built-in reports which are provided in PDF, Excel and HTML formats can be used to keep track of inventory, utilization and dependencies. Connection inventories based on connection and host type can be used to find under and over utilized accounts, applications and servers.

## ESG Lab Testing

ESG Lab ran a variety of built-in reports including connection, server and group inventory reports. Besides these useful reports, ESG Lab was impressed with the AMD Solver report options as shown in Figure Nine. The Smarts ADM Solver uses the reporting engine to provide the information needed to plan and track common business initiatives including application consolidations, application upgrades, application migrations, mergers and acquisitions, disaster recovery planning and compliance with Sarbanes Oxley regulations.

Figure Nine: Smarts ADM Solver

**Solver**

**Business Initiatives**

- Application Consolidation**  
Before running an application consolidation initiative, use this solution to identify applications, hosts and services that are under utilized and are excellent candidates for consolidation.
- Application Upgrade**  
Before upgrading your applications, use this solution to identify applications, hosts and services that are heavily utilized and are excellent candidates for upgrading. This will likely help you improve your application performance and service levels.
- Application Migration**  
Before migrating your applications to new vendors, versions or systems, use this solution to get a clear picture of your application architecture, their dependencies, and the demand placed on each application, host and service.
- Mergers and Acquisitions**  
As you acquire or divest business operations, use this solution to minimize the impact of acquiring or selling assets and ensure your business applications continue functioning uninterrupted.
- Disaster Recovery Planning**  
Create and automatically maintain accurate and up-to-date documentation of your disaster recovery plans. You can also use this solution to audit your disaster recovery plans, ensuring your business will continue uninterrupted.
- SOX Compliancy Audit**  
Section 404 of the Sarbanes-Oxley act requires you to document your key financial applications, amongst other things. This solution provides you with the necessary information required for Sarbanes-Oxley compliance readiness.

## Why This Matters

Business initiatives including consolidations, migrations and disaster recovery planning begin with a thorough assessment of application infrastructure and dependencies. Built-in Smarts ADM reporting capabilities provide a real-time assessment of the application infrastructure deployed in the data center today so that IT managers can react quickly to tomorrow's business initiatives.

## ESG Lab Validation Highlights

- ☑ ESG Lab confirmed that the agent-less packet inspection technology at the heart of the Smarts ADM architecture introduces no overhead on servers or network infrastructure.
- ☑ Using an appliance-based approach that connects to industry standard switch monitoring ports, ESG confirmed that deploying a Smarts ADM solution is straightforward and intuitive.
- ☑ Smarts ADM Discovery provided valuable insight into the applications deployed on 190 servers and accessed by 1,738 clients within an EMC lab.
- ☑ The Discovery dashboard summarized the number and type of operating systems and application services running within the lab as well as 822 application dependencies.
- ☑ The Smarts ADM console and reports were used to review an inventory of the servers, application and dependent application groups running within the lab.
- ☑ Optional High Definition Discovery, new to Smarts ADM version 5.0, was used to drill down further and obtain detailed infrastructure and application details. Details reviewed included a low level assessment of the J2EE and SQL Server traffic between a pair of dependent servers running a multi-tiered BEA WebLogic/Oracle application.
- ☑ Reporting options were tested including Smarts ADM Solver reports, which provide valuable planning information in support of business initiatives including mergers and acquisitions, data center consolidations and compliance with Sarbanes Oxley regulations.

## Issues to Consider

- ☑ The latest release of Smarts ADM includes two valuable capabilities that were not evaluated by ESG Lab: multi data center support and configuration management database (CMDB) integration. CMDB is part of the IT Infrastructure Library (ITIL) approach to data center management. The idea behind CMDB is to have a single repository for the automatic discovery of devices, systems, applications and their relations to one another. Sound familiar? It should, because that's the kind of functionality and repository that Smarts ADM delivers today. The issue going forward is standardization. While CMDB standardization efforts are in their infancy at this point in time, EMC is working with a variety of vendors to ensure that the open architecture of Smarts ADM can be integrated with other vendor's CMDB schemas in the short term and evolve to support standards in the future. As of the writing of this report, Smarts ADM has been integrated with the Atrium CMDB from BMC Corporation.
- ☑ With data centers deploying virtual servers in production environments, tracking the location of virtual machines and their relationship to physical server resources is a difficult task. The portability of virtual machines and the relative ease at which they can be re-configured makes resource contention a major issue. The real-time discovery and dependency capabilities of Smarts ADM would be an excellent tool for the management of virtual machines. EMC has advised ESG that virtual machine support will soon be added to a future release of Smarts ADM.
- ☑ The High Definition Discovery process requires unrestricted read-only administrator level access to application servers. Instead of handing out unrestricted administrative access, a number of data center administrators that ESG has spoken with are setting up read-only management domains for management applications like Smarts ADM. This requires some planning and effort, but is recommended for optimal security.

### **ESG Lab's View**

Smarts ADM is a very important step towards EMC's goal of extracting business value from the IT infrastructure. The ability to automate the discovery of application and map the dependencies in the underlying infrastructure is of tremendous value to data center managers, especially when migrating applications or planning for disaster avoidance. Obviously the larger your company and the greater the number of applications and geographically dispersed sites, the more value can be derived.

For those companies that have already invested in the EMC Smarts portfolio, Smarts ADM makes a great deal of sense. Together, these packages provide a path towards an abstracted view of your complex IT environment with the ability to map and manage applications deployed over an ever-changing infrastructure.

ESG believes the future of the data center will be focused on providing a transparent service utility. One in which business decisions will dictate infrastructure allocation and that the actual infrastructure will be highly abstracted. Users are only interested in interacting with the applications. They don't care why the application doesn't work – they just want it to work. Recently, there has been a lot of industry attention surrounding the Information Technology Information Library (ITIL) and Information Technology Service Management (ITSM) initiatives with the goal of establishing standards and a framework for the effective and efficient delivery of IT services. While neither of these initiatives specifically calls out technology, it is clear that a transparent service utility will require automated multi-domain software to detect and keep track of the underlying IT infrastructure. ESG believes that Smarts ADM software is a good first step towards a transparent service utility and recommends that EMC customers pay close attention to the benefits that Smarts solutions can deliver today.