THOUGHT LEADERSHIP WHITE PAPER

Four Steps to Faster, Better Application Dependency Mapping

Laying the Foundation for Effective Business Service Models

By Adam Kerrison, Principal Product Developer, BMC Software, and
Simon Woodward, Principal Product Manager, BMC Software

bmcsoftware
BUSINESS RUNS ON I.T.™

TABLE OF CONTENTS

EXECUTIVE SUMMARY

IT services are embedded in critical business processes across most organizations — from online commerce to back-office operations. Interruptions to — or degradations of — these services can result in serious business consequences. That's why it's so important for IT to understand the impact of changes and downtime on key applications and business services.

Service models play an important role in enabling IT to ensure the delivery of services at the levels required by the business. These models portray the relationships among business services and the underlying applications, IT infrastructure components, processes, roles, and service level agreements (SLAs) that define and support them. It's important to create and maintain accurate and complete models of the business services that drive your organization. Such models can help you to make better planning decisions and recover more quickly from application downtime.

An essential step in creating service models is developing application maps that show the relationships among the applications and the underlying IT infrastructure components that support services. Some of the challenges in developing these maps include dealing with the complexity of multitiered applications and continuous change in the IT infrastructure.

Once you have built the application maps, you can make them available in the configuration management database (CMDB) as the first step in defining business services. You can then add other information, such as users, policies, deployment blueprints, and service level key performance indicators (KPIs), to complete the service definitions.

This paper describes how to leverage advanced discovery and dependency mapping technology to help build business application dependency maps. It also includes examples of organizations that have used the approach and reviews the significant benefits they have achieved.

APPLICATION MAPS: ESSENTIAL COMPONENTS OF SERVICE MODELS

To ensure that IT is effectively supporting business demand, you must understand the relationships among business services and the underlying IT applications and infrastructure that support them. You also need to understand and define the policies, deployment blueprints, user relationships, service KPIs, and service level targets that are associated with these business services. This understanding gives you the insight to effectively manage the entire service lifecycle and focus on those issues that have high business impact.

Examples:

» Application A is deployed on a single server; application B is deployed in a clustered server environment.

  – If the server holding application A fails, the application is completely unavailable. By comparison, if a server supporting application B fails, the performance of the application may be affected, but the application is still available because of the server redundancy.

  – Application mapping enables IT to understand and visualize the connection between applications and servers and to understand the potential effects when a server fails.

» A router fails.

  – Knowing which applications and business services are affected by the failed router enables IT to prioritize its actions based on the business impact — e.g., when the router is a component of a mission-critical business service versus a noncritical business service — and to proactively communicate with the service users and owners.

As Figure 1 illustrates, application maps are essential components of service models. These maps portray the relationships among the business applications that support the service and the underlying IT infrastructure.

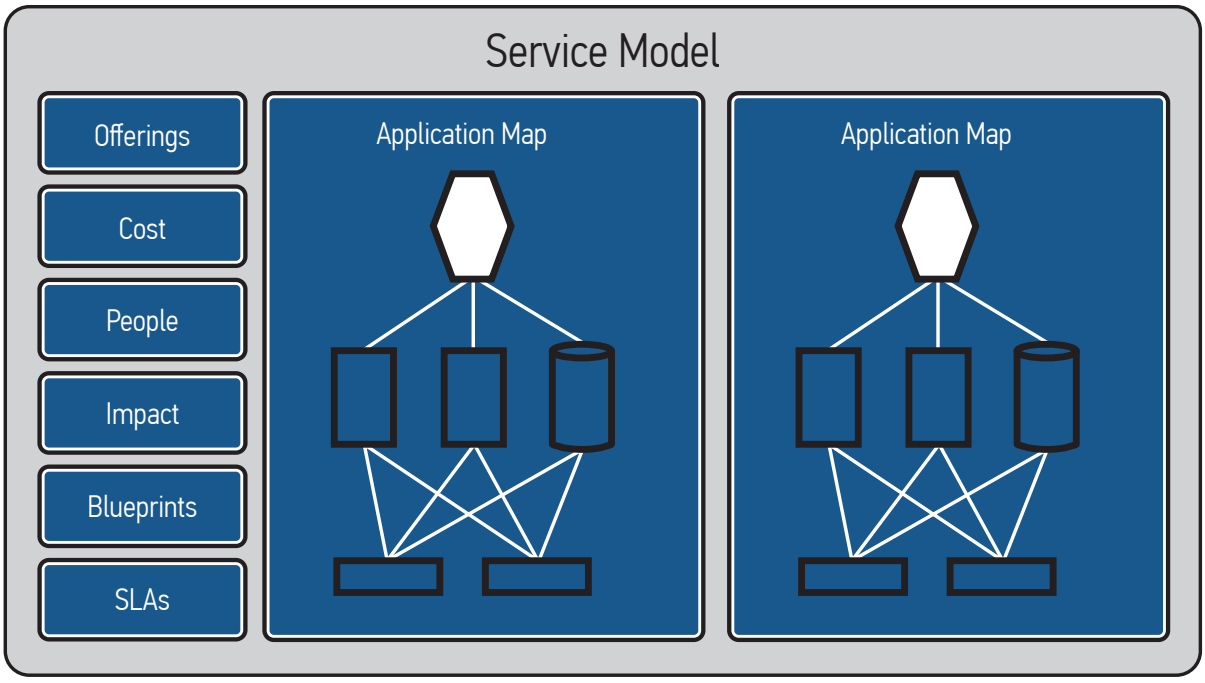

Figure 1. Service model

In most IT organizations, the people tasked with building service models, and hence application maps, usually reside in groups responsible for such activities as configuration management or enterprise architecture. In building the models, these groups need the help of application experts who are typically in application support or application architecture groups.

If yours is like many IT organizations, you may have taken a more traditional approach in creating application maps, one in which you conduct fact-finding meetings with the application support people who best know the environment. The problem, however, is that there may be a disconnect between your group and theirs. Your group needs the information that the application experts have, so you are highly motivated to obtain this information. The application experts, on the other hand, already have the data they need (often in their heads) to do their jobs, so they may not be highly motivated to expend the time and effort required to help you. As a result, getting their attention might be difficult.

Surprisingly, automation — when not done right — often exacerbates the situation. For example, automatic discovery technologies are available that gather vast amounts of data about the IT infrastructure. In the traditional approach, you might simply hand the expert a spreadsheet list of thousands of servers located by the discovery tool. You would then ask which servers support the target application and what software is involved. Providing the requested information would require the application experts to laboriously sift through reams of data — an arduous task that they may be reluctant to tackle.

BRIDGING THE GAP

So how do you bridge the gap between you and the application experts? You need a new approach to communication — one that not only analyzes and presents the discovered data in a form that motivates the experts to participate but also minimizes their time and effort.

BMC has developed a straightforward approach using the sophisticated BMC Atrium Discovery and Dependency Mapping solution. This approach leverages the company's extensive experience in helping thousands of customers worldwide to implement Business Service Management (BSM), a unified platform and comprehensive approach that helps IT organizations cut cost, reduce risk, and drive business profit.

The approach consists of four steps:

- *Step 1: Automatic discovery and dependency mapping.* Automatically create an initial map of the dependencies and relationships of the target application. (Even the most sophisticated technology cannot always create a complete picture of applications. That usually requires human intervention, which is performed in Steps 2 and 3.)
- *Step 2: Analysis and reporting.* Generate a report that presents the automatically created application dependency map in an easy-to-read format.
- *Step 3: Collaboration.* Present the report to the application expert and request the expert to validate the accuracy and completeness of the application map and to suggest changes where appropriate.
- *Step 4: Continual update.* Continually update the affected application maps each time a change is made to the IT infrastructure.

STEP 1: AUTOMATIC DISCOVERY AND DEPENDENCY MAPPING

In this step, you use automatic discovery and dependency mapping technology to generate a topology map of the application. The map depicts the physical and logical dependencies of the application components, as well as the dependencies of the components on the underlying IT infrastructure resources. This automatically generated map brings you a great distance toward a complete and accurate service model.

Creating a comprehensive and detailed map places stringent requirements on the discovery and dependency mapping technology. Today's enterprise IT infrastructures may include thousands of hardware and software components — clients, servers, mainframes, network devices, operating systems, and multitiered enterprise applications with middleware. Virtualization has increased the complexity, with thousands of virtual resources running on the physical resources. Cloud computing drives this complexity to an

even higher level. The discovery technology, therefore, must have considerable breadth and depth of coverage to ensure that it can see through this complexity to discover and gather details about all the components in the infrastructure.

However, direct discovery of the components and related details is only part of the requirement. The technology must also be capable of analyzing the directly discovered data to create a comprehensive picture of the relationships of applications to each other and to the underlying IT infrastructure.

Mapping technology is available that analyzes the discovery data using reasoning in much the same way that an application administrator would perform this analysis. For example, BMC Atrium Discovery and Dependency Mapping analyzes the communication among systems, such as commands issued to and data returned by target systems. Through this analysis, the solution makes reliable inferences about the IT infrastructure, such as identifying hosts, subnets, software instances (SIs), and business application instances (BAIs).

Because of the complexity of today's multitier enterprise applications, the discovery and dependency mapping technology must have built-in knowledge of these applications, such as how their various components fit together. For example, the technology might have knowledge of commercial enterprise applications, such as SAP®, Oracle® applications, and Microsoft Exchange®. This knowledge would include static information about the applications and publishers, as well as hardware and software version information.

Figure 2 shows an example of the knowledge that a solution might have of an instance of the Oracle E-Business Suite Field Sales application. The figure shows instances of Oracle Field Sales business applications, depicting the relevant field sales software instance, including its four modules, the E-Business Suite software instance, and the three-tier software instances: application server, Web server, and database server.

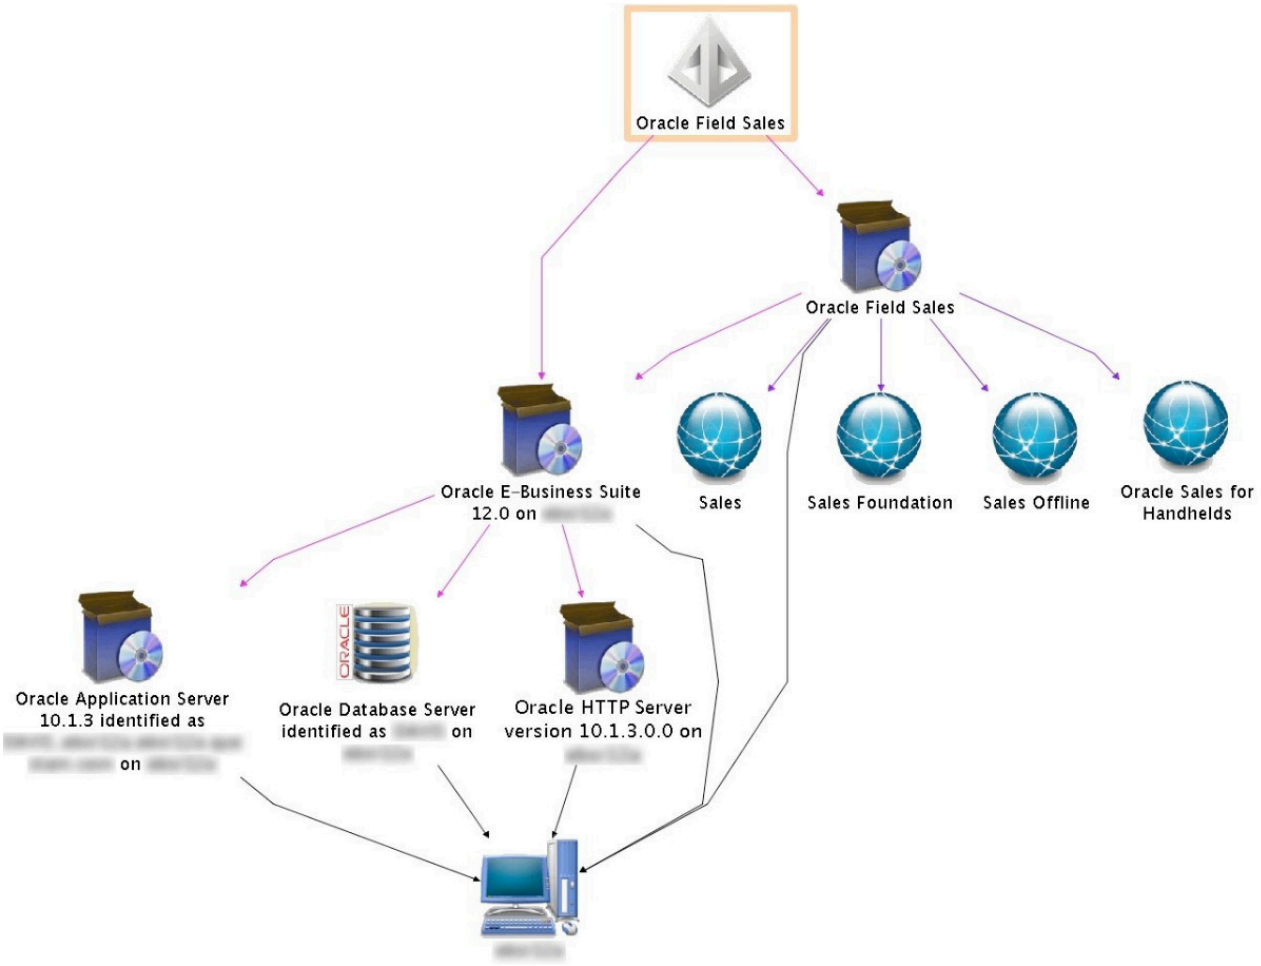

Figure 2. Map of Oracle field sales application (Identifying information has been removed.)

The technology must permit users to update this knowledge base of commercial applications as needed, such as when a new version is released — automatically, if possible. It must also enable users to easily add knowledge of custom applications.

STEP 2: ANALYSIS AND REPORTING

In this step, you leverage the analysis automatically created by the solution in Step 1 to create a meaningful document for collaboration with application experts. The solution should provide analysis and reporting tools to help you create the document.

The analysis created in Step 1 divides the discovered hosts into groups and shows the interdependencies among those groups. (See Figure 3.) Each group consists of two or more highly connected hosts.

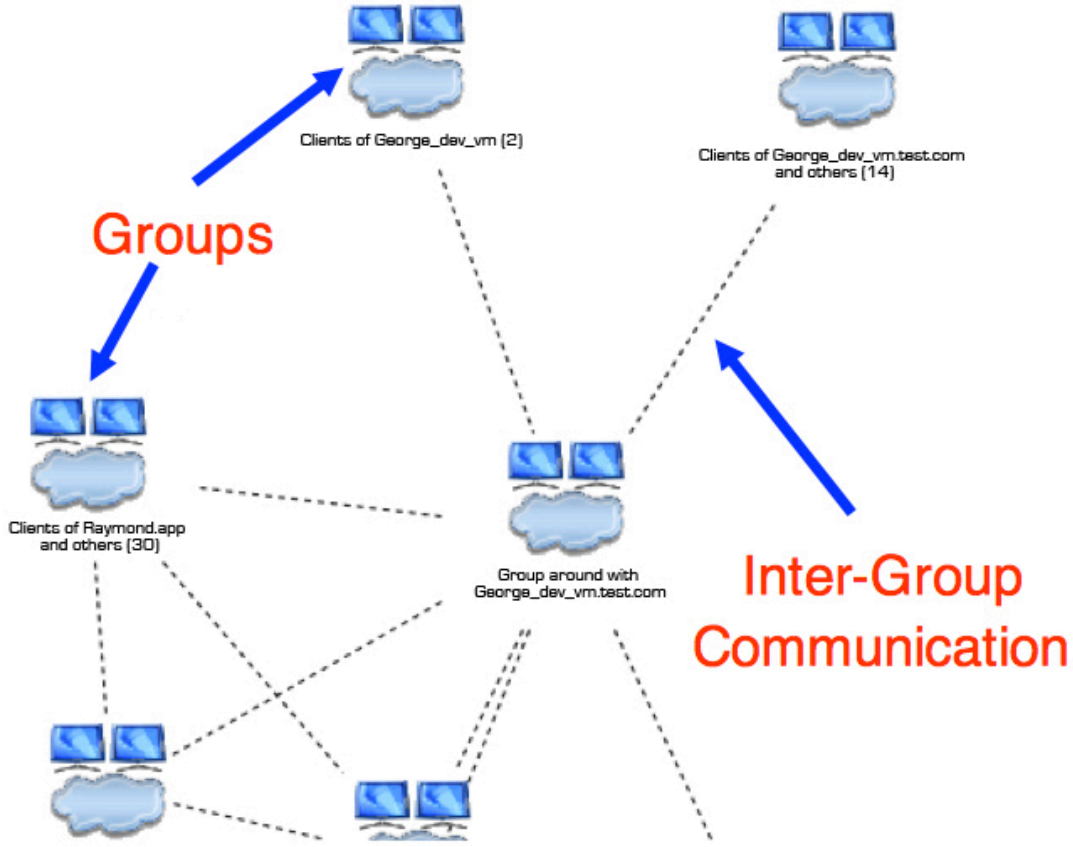

Figure 3. Hosts and their interdependencies (This view may also show the "No Relevant Communication" and "Excluded Hosts" groups so they can be accessed, although they will have no connections to the rest of the diagram.)

You can use analytics tools provided by the solution to zero in on the desired group, drill down into that group for a more granular view, and/or remove hosts to create the document you will present to the application expert. You can also request a tabular representation of the group to gain additional detail, including information about the product, ID, manufacturer, Internet protocol (IP) address, and other related information.

You then take a snapshot of the model and create a report for use in collaborating with the application expert. Figure 4 and Table 1 show a sample report of a model based on a particular host, presenting both a graphical and a tabular representation of the model.

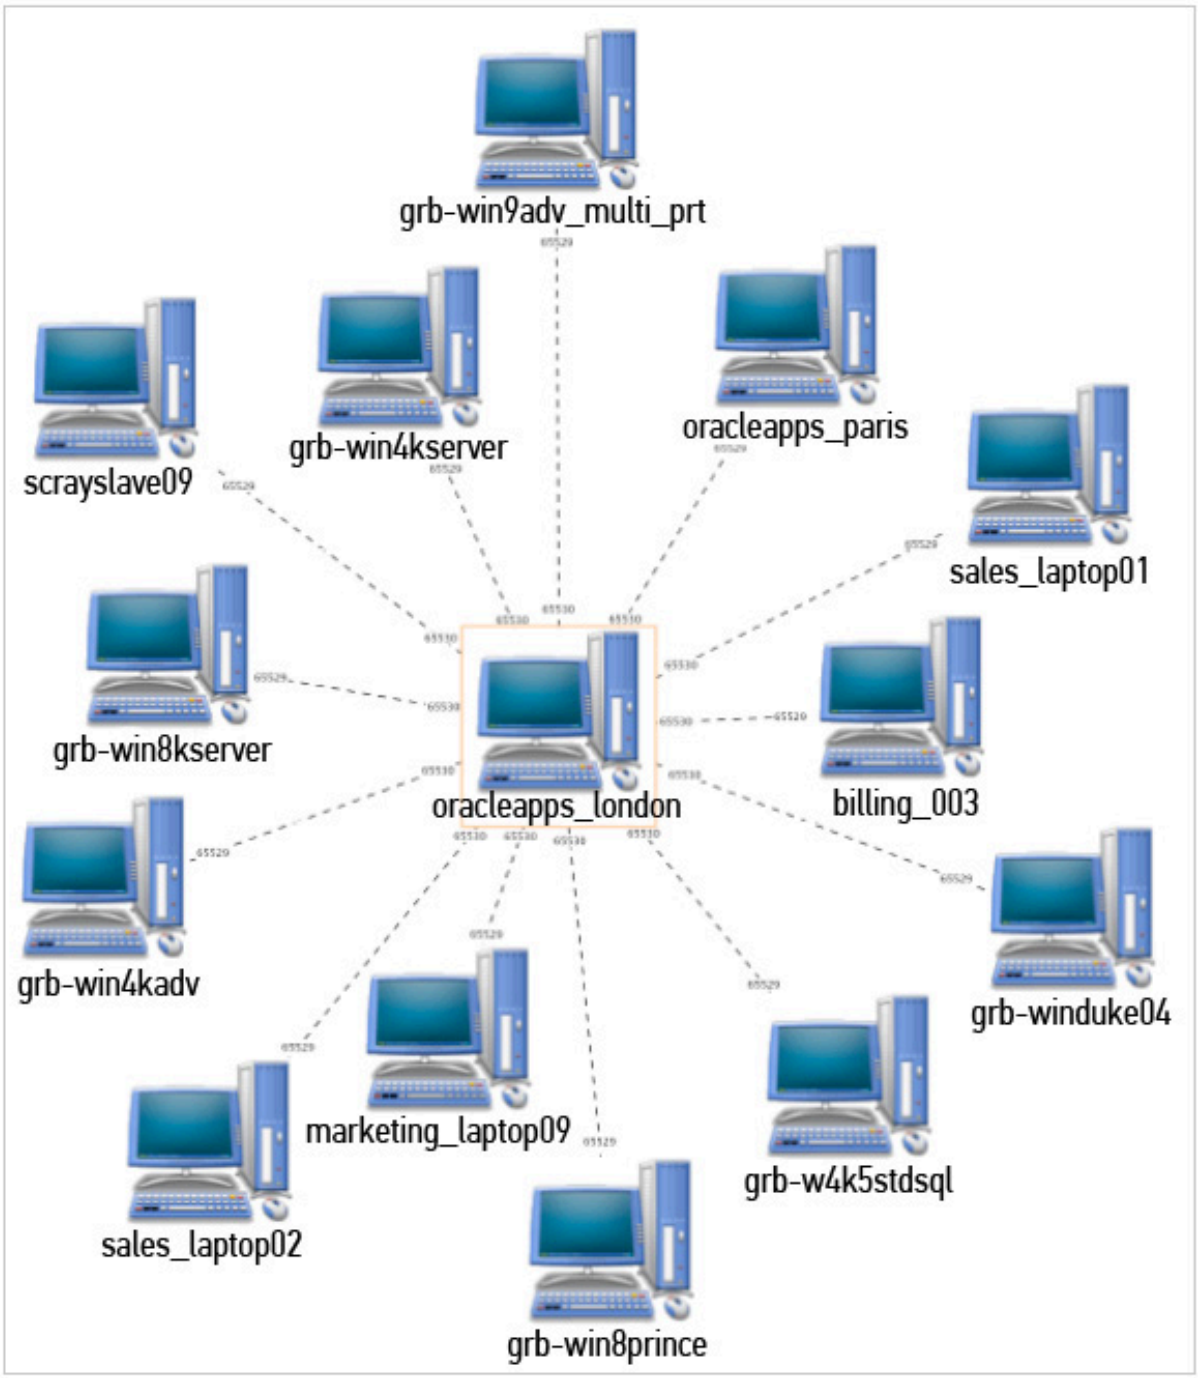

Figure 4. A host and its connections

Communication with known hosts

Local	Known Remote
183.55.71.45:65540 st harness test service.exe These are the command line arguments for the samson test...	(2 IP addresses):65529 ...st harness test service.exe These are the command line arguments for... sales_laptop02, scrayslave09
183.55.71.45:65540 st harness test service.exe These are the command line arguments for the samson test...	183.55.71.160:65529 ...st harness test service exe grb-w4k5stdsql
183.55.71.45:65540 st harness test service.exe These are the command line arguments for the samson test...	(9 IP addresses):65529 ...st harness test service.exe sales_laptop01, marketing_laptop09, grb-win8prince, billing_003, grb-win4kadv, grb-win8kserver, grb-win4kserver, grb-winduke04, grb-win9adv_multi_prt
183.55.71.45:65540 st harness test service.exe These are the command line arguments for the samson test...	183.55.71.93:65529 oracleapps_paris

Communication with unknown hosts

Local	Known Remote
151.89.35.16:1062 ...NDOWS\system32\winlogon.exe	151.89.87.37:389
151.89.35.16:23	151.89.35.819:53801
151.89.35.16:2358	185.25.91.74:80
151.89.35.16:2359 C:\WINDOWS\system32\lsass.exe	151.89.87.37:135
151.89.35.16:2360 C:\WINDOWS\system32\lsass.exe	151.89.87.37:1026
151.89.35.16:445 System	151.89.87.85:4305
151.89.35.16:65530 ...st harness test service.exe These are the command line arguments for the samson test...	183.55.71.161:65529

Table 1. Details in a tabular format

Now, you can approach the application expert with a fairly complete rendition of the application map and ask that person to validate and add to the knowledge that you already have. Compare this to the traditional approach, in which you might present a spreadsheet that contains all the infrastructure components and ask the expert to point out the relevant components and describe their dependencies.

STEP 3: COLLABORATION

In this step, you use the report created in Step 2 to collaborate with the application expert in validating and refining the application map that the discovery and dependency mapping solution has automatically built. Here, you ask the expert to review the application map and point out any inconsistencies.

The comprehensiveness of the application map and the professional quality of the report will help motivate the expert to "fill in the blanks." The report will engage the application owner and show the real value early on. The report may even reveal something of which the application expert was not aware.

Here are some examples of specific questions you might ask:

> "All of these instances of WebLogic have an Enterprise Archive (EAR) file running called stockcontrol.ear. Should I search for more WebLogic instances in your estate running that EAR?"
>
> "All of these instances of WebLogic are configured to run in a cluster. Should I include all members of that cluster?"
>
> "A few of those WebLogic instances are configured to talk to an Oracle database running over there. Should I include that database? If they start talking to a different database running a schema of the same name, should I automatically switch the links?"

"A cluster of Web servers over here is making calls into your cluster of WebLogic servers. Should I include them?"

You then use tools provided by the discovery and dependency mapping solution to make any modifications to the map suggested by the expert, such as adding or deleting hosts in a group.

STEP 4: CONTINUAL UPDATE

A major characteristic of enterprise data centers is continuous change. Consequently, the approach is based on dynamic rather than static models. That means the discovery and dependency mapping solution must be capable of not only creating dynamic models but also keeping the models updated as changes occur in the environment. Updating should be automated wherever possible to ensure that the models are current. For example, if a physical server is replaced with a more powerful one, automatic discovery would detect the change and update the server information in the model.

In some cases, such as when the current version of a commercial enterprise application is updated to a new version that includes a new application component, it may be necessary to update the knowledge base to reflect the addition. The solution would then include the new component in subsequent environment scans.

REAL-WORLD EXAMPLES

This section describes two organizations that have implemented BMC Atrium Discovery and Dependency Mapping and the benefits they have realized.

GLOBAL BUSINESS CONSULTING FIRM

This firm provides consulting services that contribute to the business transformation and economic performance of organizations based on in-depth knowledge of client industries and processes. The firm took on two outsourcing assignments, each from a large government organization. Each assignment involves managing the client's server estate and optimizing the IT environment.

In both assignments, the company found itself in a situation familiar to most outsourcing companies in that the server estate was poorly documented. Consequently, the firm had to quickly create an accurate model of each estate that included the infrastructure dependencies of the business services. The company created the models using BMC Atrium Discovery and Dependency Mapping. The results have been gratifying:

Assignment 1

- Accurately audited an IT estate of 1,600 servers and model dependencies of 76 business services, all within nine weeks
- Achieved 92 percent coverage of the server estate and had the data validated by the client as being 100 percent complete and accurate
- Reduced the number of resources required to maintain the CMDB data by 80 percent

Assignment 2

- Discovered 868 servers against a record of 350 servers
- Completed initial discovery scan in four hours, with detailed scan completed within two days
- Reduced costs by GBP 160 million over the term of the contract through elimination of manual audits
- Improved currency of CMDB data from 60 days to daily refresh cycle

LARGE EUROPEAN MORTGAGE AND SAVINGS PROVIDER

The IT organization of a mortgage and savings provider with more than 13 million customers wanted to gain measureable cost savings through virtualization and server consolidation. Its goal was to virtualize 80 percent of the production Wintel servers. It recognized that the business benefits could be achieved only if it was possible to mitigate the risks. This required a clear understanding of all servers, applications, and critical interdependencies within their IT environment. A manual audit of the IT environment was deemed too expensive and inadequate to deliver the deep, accurate, and actionable intelligence required, particularly around application dependencies.

The IT organization addressed the challenge using BMC Atrium Discovery and Dependency Mapping. Here are the results:

» Identified 76 percent of their physical servers that could be virtualized
» Analyzed 654 Wintel servers and modeled 189 business services in four weeks

CONCLUSION

This paper has presented a pragmatic, four-step approach to building application maps. The approach, based on technology found in BMC Atrium Discovery and Dependency Mapping, bridges the disconnect that has existed between those people tasked with application mapping and the application experts. By removing a major obstacle to collaboration between these two groups, the approach greatly facilitates the creation and maintenance of application maps.

Using the approach presented in this paper, IT organizations can build the application maps they need quickly and inexpensively. They can then leverage these maps to build comprehensive and accurate service models that enable effective business service lifecycle management and increase IT value contribution to the business.

For more information on BMC Atrium Discovery and Dependency Mapping, visit **www.bmc.com/products/product-listing/BMC-Atrium-Discovery-and-Dependency-Mapping.html**.

Business Runs on IT. IT Runs on BMC Software.

Business thrives when IT runs smarter, faster and stronger. That's why the most demanding IT organizations in the world rely on BMC Software across distributed, mainframe, virtual and cloud environments. Recognized as the leader in Business Service Management, BMC offers a comprehensive approach and unified platform that helps IT organizations cut cost, reduce risk and drive business profit. For the four fiscal quarters ended March 31, 2011, BMC revenue was approximately $2.1 billion. Visit **www.bmc.com** for more information.

About the Authors

Adam Kerrison, principal product developer at BMC Software, has 20 years of experience in delivering software products into operational IT environments, giving true and rapid return on investment. He cofounded Tideway Systems in November 2002 and drove the technical direction of the company's Foundation product until its acquisition by BMC Software in October 2009. Kerrison was previously CTO at Micromuse (acquired by IBM) with responsibility for direction and roadmap of Netcool products and a major architectural change in the product suite. At Micromuse, he worked his way through the ranks from Support Engineer to Senior Vice President of Development to CTO. Between 1989 and 1990, Kerrison worked as a software engineer for GEC Avionics on real-time modules for jet fighter head-up displays, relational database design, and implementation. He has a Bachelor of Science degree in Computer Science from Greenwich University, United Kingdom.

Simon Woodward, principal product manager for BMC Atrium Discovery and Dependency Mapping at BMC Software, has 15 years of experience in product management and development of discovery solutions. He was the product manager for Tideway Systems prior to the company being acquired by BMC and previously spent many years with companies such as Hewlett-Packard and Platinum Technologies, working his way up from an engineering role through architecture and development management positions before moving into product management.