

Achieve Cloud Migration Success

Deliver faster migration across the lifecycle with BMC Discovery, dependency mapping, and cost analysis



Table of Contents

1 EXECUTIVE SUMMARY

2 NOT IF, BUT WHEN: MAPPING THE MOVE TO THE CLOUD

FIVE PHASES OF A SUCCESSFUL CLOUD MIGRATION

3 SPOTLIGHT – PHASE 2: PORTFOLIO DISCOVERY AND PLANNING

TYPICAL CHALLENGES OF PORTFOLIO DISCOVERY

People

Processes

Tools

4 ANATOMY OF AN ENTERPRISE-GRADE DISCOVERY SOLUTION

5 HOW TO CHOOSE THE RIGHT DISCOVERY TOOL FOR YOUR CLOUD MIGRATION

Step 1: Inventory

Step 2: Dependency Mapping

Step 3: Grouping

Rehost

Refactor/Re-architect

Rebuild

Replace/Repurchase

Retain

Retire

Steps 4 to 6

Step 5: Cost Analysis

Post-migration

8 BMC SOLUTIONS SUPPORT CLOUD MIGRATIONS AT SCALE

What Makes Discovery Different?

Manage Ongoing Costs with TrueSight Cloud Cost Control

Case Study: A Successful Cloud Migration

9 CONCLUSION

Executive Summary

Moving to the cloud is a reality for almost every business today. Whether a single application or full-scale migration, organizations of all sizes are benefiting from cloud computing.

The process of migrating your applications and services to a cloud environment, however, can be more complex than it seems. A lack of understanding of its phases and required tools can pose a significant barrier to success, leading to schedule slips, budget overruns, compliance violations, and other disruptions.

A successful migration is defined by five distinct phases. In this white paper, we'll examine the **portfolio discovery and planning phase**. This phase provides IT teams with the visibility and insight necessary to identify and migrate the right components at the right time and in the right way. It builds the foundation for cloud migration by informing every stage of the process, from the business plan to application mapping to migration, validation, operations, and beyond.

With the right discovery strategy, you can set up your cloud environment for success through every step of its lifecycle—before, during, and after your migration project.



NOT IF, BUT WHEN: MAPPING THE MOVE TO THE CLOUD

Cloud migration needs to be treated as a top-down, company-wide, and cross-functional initiative with cross-team collaboration. It is likely that at your organization the cloud migration process has already begun, from the line of business bringing in SaaS apps to IT testing the waters for larger-scale migration projects. If so, you're not alone.

Statistics from multiple industries demonstrate that the move to the cloud is irreversible, fully in flight, and with no signs of slowing down:

- Cloud computing spend has been growing at 4.5 times the rate of IT spending since 2009, and is expected to be 6 times more than that of IT spending from 2015 through 2020.
- Worldwide spending on public cloud computing will increase from \$67B in 2015 to \$162B in 2020, attaining a 19% CAGR.
- By 2018, at least half of IT spending will be cloud-based, reaching 60% of all IT infrastructure, and will be 60–70% of all software, services, and technology spending by 2020.

The momentum of the rapid shift to cloud computing is obvious, but the “how” isn't as straightforward. While a startup may enjoy the luxury of creating a cloud environment from scratch, most enterprises face a more daunting brownfield environment laden with legacy applications and a large, complex footprint. For these organizations, migration isn't as simple as spinning up an instance or just “ripping and replacing.” It requires the right strategy and tools to get it done successfully.

FIVE PHASES OF A SUCCESSFUL CLOUD MIGRATION

Migrating to the cloud on time, on budget, and with the desired results doesn't happen automatically. It requires a comprehensive, multi-phase plan that includes checkpoints for operational adjustments. Additionally, this type of initiative typically falls within the scope of a larger transformation effort rather than occurring in isolation, making careful planning and strategic execution especially important.

Large-scale migrations that meet or exceed expectations commonly follow a five-phase approach:



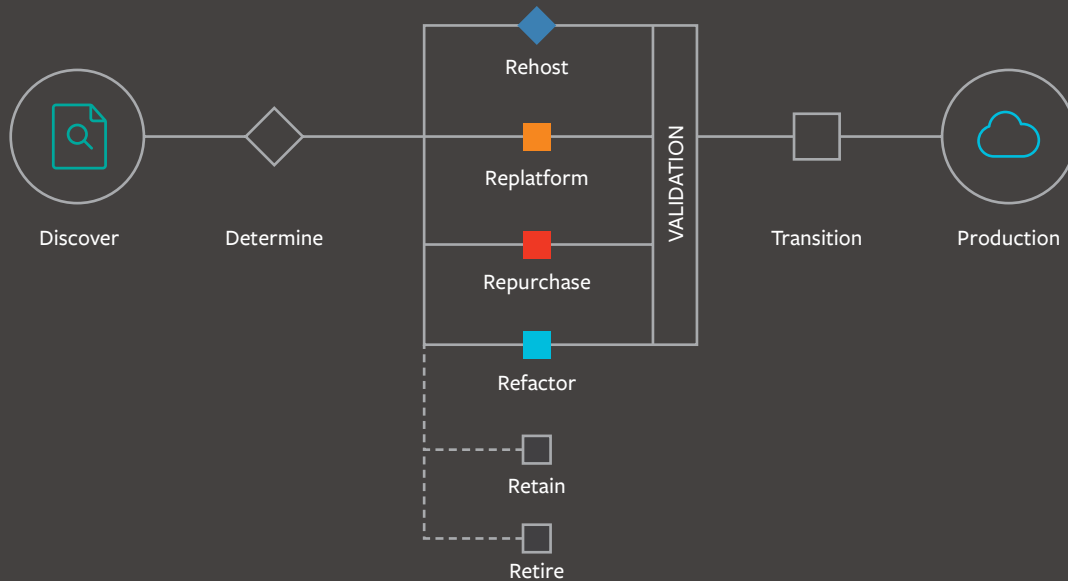
- **Phase 1: Migration Preparation and Business Planning**
Crystalize your objectives and build your business case. Factor in the age and architecture of your existing applications and constraints. Review any “lessons learned” from earlier forays into the cloud.
- **Phase 2: Portfolio Discovery and Planning**
Gain an in-depth understanding of your IT portfolio, the dependencies between applications, and based on this knowledge, the type of migration strategies that you will need to meet your business objectives.
- **Phases 3 and 4: Designing, Migrating, and Validating Applications**
Shift from the portfolio level to designing, migrating, and validating individual applications. Start with your least complicated use cases to build foundational knowledge and organizational support.
- **Phase 5: Operate**
Iterate on your new foundation, turn off old systems, and refine your new operating model. Review and revise your plans around people, processes, and technology to create a migration model that constantly improves.

SPOTLIGHT – PHASE 2: PORTFOLIO DISCOVERY AND PLANNING

Each of the stages defined above plays a critical role in the success of a cloud migration project. Let’s take a closer look at phase 2: portfolio discovery and planning.

The need for thorough discovery and dependency mapping is no surprise to IT teams tasked with the shift to cloud; you need to know what you have in order to move it effectively. A comprehensive discovery phase influences the process up and down the ladder and sets the tone for the migration. It feeds back into phase 1 of the migration process, providing the information required to build a data-driven business case and migration plan, as well as feeding into the subsequent phases—informing the approach for how to migrate and operate your infrastructure most efficiently at the portfolio and application levels. Cost visibility and modeling are also critical, allowing alignment with the commercial model of public cloud providers like AWS.

In its latter role, sophisticated discovery data allows IT to make informed decisions about the most appropriate strategy for each application, such as rehosting (a.k.a. lift-and-shift), re-platforming (e.g., switching middleware), repurchasing (e.g., using PaaS or SaaS instead), refactoring, retaining on-premises, or even retiring. With in-depth knowledge about IT assets across your infrastructure—from the data center to public, private, and hybrid clouds—you can better decide among your various migration options.



Six Common Migration Strategies: Condensed View Diagram

Finally, with proper documentation of the source environment including detailed dependencies, architects and their teams gain the data they need to design the cloud-based applications to help operations teams sequence the migration.

TYPICAL CHALLENGES OF PORTFOLIO DISCOVERY

Although portfolio discovery is a critical element of successful cloud migration, this capability remains underdeveloped in many organizations. The resulting difficulties fall into three broad areas: people, processes, and tools.

People

A shortage of cloud skill sets can have a major impact on migration success. Given the fast pace and competing demands of modern IT, it can be hard to pin down subject-matter experts to perform in-depth analysis of existing application deployments. When such analyses are performed, the manual processes typically used often result in errors such as incorrect or incomplete output.

Application owners and infrastructure teams work in silos, leading to gaps in skills and misunderstandings. In the absence of a more formal, rational process, portfolio knowledge tends to be tribal rather than institutional, making it vulnerable to loss through staff turnover.

Processes

Migration goals are often not correctly documented, for example, in terms of the scale of change, data quality, etc.—making it difficult to guide, measure, and assess the success of the effort. The lack of an application portfolio can make it impossible to establish a coherent overarching strategy, whether for mass migration or for an application-by-application or defined-application-batch approach. Without establishing each application’s dependencies with other processes, organizations run into challenges with software licensing (license models and ownership for migrated applications), compliance (requirements may be similar or different as applications move to the cloud), and service management (change windows and procedures, incident response, etc.).

Tools


Inventory tools generally fail to capture or accurately represent dependency information, especially when they rely only on observed communications, and are unable to reflect the business context of infrastructure components. Given the fast pace of change in modern IT environments, these tools are typically neither scalable enough nor fast enough to guarantee the accuracy of their data at the moment of migration. Solutions limited to point-in-time discovery are unable to fulfill the requirements for ongoing post-migration management.

Agent-based solutions are heavier to deploy and their knowledge is limited to the component on which they are deployed, though some organizations favor this approach because they feel more secure running everything in their own environment before opening up to the right cloud platform. On a fundamental level, to leverage the information and investments already present in the organization as well as to enrich the data available to optimize migration, a discovery tool must integrate or at least co-exist well with existing application performance management (APM) and configuration management database (CMDB) tools in the environment.

ANATOMY OF AN ENTERPRISE-GRADE DISCOVERY SOLUTION

The right tool can help you avoid these pitfalls and actively contribute to the success of your cloud migration. What does that look like for your migration project’s phase 2: portfolio discovery and planning?

As the name of the phase suggests, this step in the process can actually be broken down into two parts: discovery and planning.



Discovery solutions play in the first half of the process. The most basic solutions begin and end by taking an inventory of your IT assets. More sophisticated tools continue by identifying dependencies and grouping components together via application maps to understand the various “buckets” within your infrastructure.

Planning, the second half of the process, requires an understanding how your applications perform and what they cost in order to provide a recommendation on what should be migrated and how.

HOW TO CHOOSE THE RIGHT DISCOVERY TOOL FOR YOUR CLOUD MIGRATION

Let's look at the required functionality for each of the three steps within the discovery portion of phase 2. In other words, what does your discovery tool need to do to successfully support your cloud migration?

Step 1: Inventory

The inventory aspect of discovery and dependency mapping is theoretically the “easy part”—except even at this stage, many tools lack the breadth and depth to provide the visibility you need. To achieve a comprehensive and accurate inventory of your IT infrastructure, your discovery solution must cover:

- Servers: physical and virtual, hypervisor, OS, CPU, RAM, disk
- Software: all software assets, including end of life, plus databases and websites
- Network devices: switches, load balancers, etc.
- Storage: devices and their logical partitioning

A complete application inventory is essential as well, of course. Because applications in the cloud can choose to use elastic resources that auto-scale, mapping performance objectives to resources and their utilization makes it possible to optimize resources while meeting performance requirements. For this reason, utilization metrics for each of the components listed above must be understood on an application-by-application basis.

A complete inventory sets up the scope of what needs to be understood and managed across your migration project. But it's important not to stop there.

Step 2: Dependency Mapping

To fuel your cloud migration project, you need to understand not only what you have, but how it all works together. Documenting these relationships enables you to identify migration sequences, ensure minimal downtime, guarantee comprehensive test plans, and map redundancy and availability, among other benefits.

Key dependency data provided by your discovery solution should include:

- Software dependencies—e.g., web, application, and database tiers, and clustered software configurations
- Containers and microservices
- Server-to-storage relationships to understand how the data flows
- Hybrid application deployments
- Host-to-edge network relationships
- Hardware and software load balancing
- Disaster recovery setups

This list may be longer depending on the complexity of your infrastructure. Dependency data at scale is particularly critical at the enterprise level, where clarity into the many interconnections can make or break the success of your cloud migration.

Step 3: Grouping

Now that you know what you have and how it works together, you need to be able to visualize and consume that information. This is where application mapping comes in. Your discovery solution should be able to easily group components into application maps or models that articulate dependencies, so you can:

- Ensure that business impact is duly assessed
- Support business continuity objectives
- Make more informed decisions continually as to what can and cannot be migrated

Grouping also helps with the essential process of evaluating the workloads and applications that are candidates for moving to the public cloud. When assessing the readiness for running applications in the public cloud, all the options need to be considered, from lift-and-shift to re-architecting the solution into a more cloud-friendly form. The considerations must be measured against the goals of your cloud strategy and potential business risk incurred. The key is to examine all applications and workloads under a consistent framework and not on an ad hoc basis.

Rehost

With this lift-and-shift approach, an application is redeployed to a cloud-based platform without modification of its code. This can be a good way to achieve rapid scale to meet business needs. However, established applications that were not designed for efficient use of infrastructure will most likely cost more to run in a public cloud. Because of this, a simulated migration is recommended for rehosting applications to prevent cost surprises.

Refactor/Re-architect

Refactoring/re-architecting involves modifying the application, application framework, or runtime environment. This can include making application code or configuration changes to attain a tangible benefit from cloud migration without making major changes to the core architecture of the applications—for example, swapping out a database server for a cloud service equivalent to reduce the amount of time you spend managing database instances.

Rebuild

For applications written in-house, redesigning and rebuilding a cloud-native application on a provider's PaaS may be worth the investment. This can be the right choice for applications that are business-critical, but not designed to take advantage of the services offered on a cloud platform. In addition, non x86-based applications, such as mainframe and midrange applications that rely on operating systems other than Linux and Windows, will need to be rewritten. This is the most expensive option, but the investment to rewrite an application may be worthwhile if your goal is to boost agility, improve performance, reduce costs, and improve business continuity.

Replace/Repurchase

For commercial, on-premises applications, replacement with a SaaS version from the same vendor may be the best solution. Even if the preference is to run on-premises, many ISVs have upgraded their applications to run better on cloud platforms, and achieving your goals could be a matter of upgrading the application to a more current version.

Retain

Sometimes it makes good business sense to keep some applications on-premises for reasons of cost, security, or compliance. Additionally, not every application can benefit from a cloud platform; those with static workloads and no need for agility, and currently running on stable systems, are good candidates to retain on-premises.

Retire

The application evaluation phase also provides an opportunity to identify applications and workloads that are no longer needed or lack the business justification to warrant the ongoing cost to support them. This is a great time to rationalize your portfolio.

Grouping is not a one-size-fits-all activity. Your tool should be able to derive an application model from any piece of information about that component for optimal flexibility and time savings. It should also be customized to your organization with application maps built for business value. This includes the ability to define rules, and to query and visualize the discovery and dependency data in the way that best enables your stakeholders to understand and act on the information.

Steps 4 to 6

While discovery solutions don't explicitly address the second half of the portfolio discovery and planning phase (planning), they ideally feed data into steps 4 to 6 so that other solutions can more effectively deliver their analysis and recommendations.

The number-one requirement here is the mature use of exports and APIs. Once you've grouped your applications and gained visibility into their dependencies, you can export that information and overlay it with cost and performance data. This adds another layer to your analysis for deeper insight into the overall state of your environment.

Step 5: Cost Analysis

Although beyond the scope of discovery solutions, cost analysis is a critical element of cloud migration planning and merits a brief exploration here.

When migrating applications to public cloud platforms, a financial assessment can be just as important as technical assessment. This is especially true if the migration includes rehosted or refactored applications; without understanding the cost difference for running these applications on-premises vs. in a public cloud, you're likely to get an unpleasant surprise. Even so, in a [recent BMC survey](#), 40 percent of respondents stated they are unclear of their costs associated with cloud—even though lower cost was the primary driver (45 percent) for moving to public cloud.

Every IT organization should have a history of resource usage and workload patterns for their applications, along with the unit cost of infrastructure resources. Many organizations gather this information routinely and use it for chargeback or showback for IT infrastructure costs. Absent this information, you can still gain some level of insight based on, at minimum, a unit cost for on-premises compute, storage, and network resources for the workloads and applications you want to migrate.

Determining AWS resources for migrating applications and workloads is easily done with an automated solution such as **TrueSight Cloud Cost Control**. Based on the historical resource utilization of the infrastructure resources associated with the application, the best fit for type and size of compute services is recommended. A cost comparison for on-demand and reserved instances helps you make an informed decision based on your knowledge of the life of the application and the available pricing options.

If applications or workloads are run periodically or have a short life, such as an application test period, on-demand pricing may be the lower-cost option even though the monthly rate is higher. However, for applications or workloads that are an integral part of your business and run routinely, reserved instances are a better cost choice.

This analysis can also be performed manually, though this can take weeks of work depending on the number of applications and workloads you are analyzing. An automated solution is much preferable to ensure fast, accurate, and comprehensive analysis.

Once the analysis is completed, you should view it as a snapshot in time. IT environments are naturally dynamic, so costs will change over time. An ongoing cost management practice is essential to regulating operating and capital expense for a hybrid or cloud environment. A good cost management practice requires automation that includes:

- **Predictive analytics for budget management** – Identifying anomalous behavior that can lead to budget overages helps control costs and avoid surprises.
- **Forecasting for budget planning** – Visibility into monthly, quarterly, and annual costs based on historical trends helps set a baseline. In addition to capacity planning for future needs (both onboarding and offboarding workloads and applications), this provides the insight for budget planning.
- **Flexible views and reporting** – Information about costs and resources should be made easily available to stakeholders in the context of their business needs.
- **Analysis of reserved instances** – By ensuring that you are maximizing your committed spend, you can increase the value of your cloud investment.
- **Identifying oversized resources** – Savings can be achieved by identifying idle VMs, unused or underused storage and compute resources, and lower-cost regions.

With the right tools and analysis, you can achieve your goals of greater agility and lower costs by extending your IT environment to include the public cloud.

Post-migration

Finally, consider your post-migration needs when selecting a discovery solution. Discovery tools not only provide valuable insight before and during a move to the cloud, they also play a key role in iteratively optimizing and improving your ongoing operations.

Consistent, clear visibility is especially imperative in today's multi-cloud environments. As evidenced by the six strategies for cloud migration listed earlier, you can take different approaches for different applications and services. The completion of all cloud migration projects does not result in everything-in-the-cloud; some components may remain in the data center, some may shift to a hybrid approach, and even those workloads that do move to cloud will do so in a variety of ways.

Comprehensive discovery and dependency mapping ensures that the new way a service or application is being delivered is optimal. It benefits operations teams by providing always up-to-date application model documentation as well as application support around incident, problem, and change management. It improves service-aware monitoring, capacity optimization, and automation capabilities.

Finally, thorough, ongoing discovery and dependency mapping can be used as a source of information to validate efforts and results across IT, such as:

- **Developing security models** – network access, component versions, subnets, etc.
- **Creating performance and availability models** – load-balancing, elasticity, resilience, monitoring
- **Ongoing management** – configuration management, continuous delivery, etc.

For security, with automated discovery and dependency mapping, you can formulate and execute a plan to resolve known vulnerabilities with the least amount of risk. According to WhiteHat Security, 80 percent of attacks go after a known vulnerability. Discovery solutions not only ensure that you're aware of all risks, but that you're armed with the information you need to address them.

360-degree visibility enables faster response to security threats as well. A comprehensive discovery tool helps you evaluate business impact more quickly and conclusively, then respond to critical events with reduced risk of incorrect assignment and assessment. That means you can proactively identify trends, abnormalities, and security vulnerabilities, and automatically remediate them before they impact the business.

BMC SOLUTIONS SUPPORT CLOUD MIGRATIONS AT SCALE

One discovery and dependency mapping tool meets all the criteria necessary to support the scale and depth of an enterprise cloud migration project, such as the move to AWS: **BMC Discovery**.

BMC Discovery automates asset discovery and dependency mapping to better gain visibility into management of digital services in cloud and on-premises environments. It provides the trusted foundation for building application maps that enable digital transformation to be more aware, secure, and cost-transparent.

Each scan delves into the information and dependencies for all software, hardware, network, storage, and cloud services, providing IT with the context needed to create an application map from any piece of information about it. This allows cloud migration projects to rely on robust inventory data, dependency maps, and automatically updated component groupings.

Once your applications are migrated, BMC Discovery enables streamlined operations by offering multi-cloud dependency maps to ITSM, ITOM, and SecOps processes. It models all of your dependencies in minutes, serves them up from a single pane of glass, and offers **an average of 470 percent ROI in the first five years, including a typical payback period of just eight months.**

What Makes Discovery Different?

BMC Discovery addresses both comprehensive IT asset discovery as well as application dependency mapping. For 15 years, the solution has been providing the broadest data center and cloud coverage, currently detecting over 4,000 types of assets and 1,000 types of relationships. Its Technology Knowledge Update (TKU) content is updated monthly, providing extensive documentation of every supported technology, including mainframe, servers, storage systems, network devices, software, hardware, and cloud services.

This out-of-the-box content enables Start Anywhere Application Mapping, a unique approach to automatically grouping and documenting application models. Whether your environment is heavily clustered, highly segmented, or has shared DB, web, and middleware software, or applications spanning multiple AWS regions, we can understand it and visualize it.

Comprehensive discovery also enables blind spot detection, which is critical for identifying vulnerable systems and meeting regulatory compliance requirements. BMC Discovery meets the highest levels of security, providing an agentless, virtual-hardened appliance architecture that is STIG compliant and Common Criteria EAL 2+ certified.

Customers consistently rave about how easy BMC Discovery is to deploy and use. To browse customer feedback and case studies, visit our [TechValidate case study portal](#).

Manage Ongoing Costs with TrueSight Cloud Cost Control

A valuable complement to BMC Discovery, TrueSight Cloud Cost Control enables ongoing cost management by providing insight and control over capital and operating expenditures. A single view of on-premises and public cloud infrastructure expenditures allows you to track and analyze infrastructure costs and utilization, identify wasted spending, and forecast future costs. The ability to easily simulate cloud migrations and compare on-premises and public cloud infrastructure costs allows applications to run on the most cost-efficient infrastructure. An application and service view of infrastructure resources and cost makes actionable information easily accessible by business and IT stakeholders. Thus, TrueSight Cloud Cost Control helps you align cloud costs and usage with your IT business strategy.

Case Study: A Successful Cloud Migration

A multinational power company moved applications to AWS as part of a broad digital transformation designed to increase its agility. After identifying over 1,700 applications to move to AWS cloud services—roughly 80 percent of its total portfolio—the company needed to organize them into logical groups and set a priority for moving them. To do this, they needed to identify the infrastructure resources associated to the applications and understand any dependencies.

The company used BMC Discovery to identify infrastructure assets and map application dependencies. This information was used to populate the company CMDB, which was foundational for identifying application and business service infrastructure resources and dependencies to determine what would remain in the data center and what applications and infrastructure would be migrated to AWS. Applications were grouped by size as a proxy for migration complexity, including groups for SAP-related applications, a relatively small number of compute-intensive applications, and a larger proportion of less compute-intensive applications. For the most part, the company chose a lift-and-shift migration model.

As the migration began, the company realized they were moving and paying for underutilized VMs. To eliminate this wasted spend, they used TrueSight to help right-size VMs. An active dashboard was created showing each business service and the associated infrastructure resources. Utilization data was regularly collected and used to report underutilized CPU, memory, and storage. Automated recommendations were provided and with BMC Atrium Orchestrator, automated actions were taken to make the appropriate change.

The combination of AWS, BMC Discovery, BMC Atrium Orchestrator, and TrueSight has enabled an automated migration process as well as ongoing resource optimization. Applications have the infrastructure resources needed for meeting performance expectations while resources are optimized and operating budgets are met.

The outcome: the shift to AWS has yielded savings of up to 60 percent in storage costs and 20 percent in computational power, and reduced the time required to provision from three to four weeks to two days.

CONCLUSION

Large-scale cloud migration projects are either already or will become a top priority for most organizations in 2018 and beyond. Ensure that your migration project is fueled by the data and analysis necessary to make the move a success—from implementation through ongoing operations. Providing comprehensive visibility, scale, and insight to accelerate and optimize every part of the cloud lifecycle, a discovery and dependency mapping tool like BMC Discovery, complemented with cost management powered by TrueSight Cloud Cost Control, is an essential element of effective cloud migration for the digital enterprise.

FOR MORE INFORMATION

To learn more about BMC Discovery, visit bmc.com/trydiscovery

To learn about about TrueSight Cloud Cost Control, visit bmc.com/cloudcost

BMC is a global leader in innovative software solutions that enable businesses to transform into digital enterprises for the ultimate competitive advantage. Our Digital Enterprise Management solutions are designed to fast track digital business from mainframe to mobile to cloud and beyond.

BMC digital IT transforms 82 percent of the Fortune 500.



BMC, BMC Software, the BMC logo, and the BMC Software logo, and all other BMC Software product and service names are owned by BMC Software, Inc. and are registered or pending registration in the US Patent and Trademark Office or in the trademark offices of other countries. All other trademarks belong to their respective companies. © Copyright 2018 BMC Software, Inc.

