



Think fast.™

Riverbed Cascade™ Discovery for BMC Atrium

Cascade Discovery provides discovery and dependency mapping for BMC Atrium without the need for agents or scanners by automatically analyzing network flow information to identify application/server and server/server dependencies. This unique method for monitoring application usage by using the existing network infrastructure provides detailed knowledge of how every application is being used, by whom, and when on a continuous basis. Cascade Discovery's passive approach enables organizations to maintain up-to-date, as-is CMDBs with minimal need for intrusive scanning.

Overview

Your CMDB is only as effective as your ability to populate it with complete and accurate information. Identifying and maintaining all of your organization's configuration items and their dependencies is a challenge. Manual inventory methods are expensive, time consuming, and error-prone. Approaches which use scanning and agents can be expensive to deploy and can affect the performance of systems and networks. Cascade Discovery offers a new way to accurately and efficiently keep your CMDB up to date. Scanning-based approaches typically have coverage gaps due to their inability to scan the entire environment. Scanning solutions are only able to identify dependencies that exist at the time of the scan. This lack of knowledge of dependencies outside the typical scan window creates a large gap in the CMDB's effectiveness. Agent-based solutions can provide a dependency view over time but are expensive to deploy and can rarely be deployed broadly enough to provide an adequate dependency mapping solution. Cascade Discovery offers a new way to accurately and efficiently maintain up-to-date application dependency relationships in your CMDB.

As a BMC Atrium user, you have a single, consolidated view of your IT infrastructure through a federated CMDB environment. The Atrium CMDB is populated via a scanning process. While this scanning process provides deep information about your configuration items (CIs), there are a few challenges associated with this approach:

- Certain systems aren't available to scanning due to firewalls and/or network configuration or credential issues
- Scanning can only provide information on application dependencies that are running during the scan, which typically lasts less than a minute
- Pervasive scanning can affect network performance

Cascade Discovery for BMC Atrium addresses these challenges by providing a new way to passively discover configuration items and their dependencies across the entire enterprise. It complements your BMC Discovery deployment to ensure complete and accurate population of your CMDB without time or coverage gaps.

How it Works

Cascade Discovery's primary data source is network traffic flow data – such as NetFlow, sFlow, cflow, and IPFIX – provided by most routers and switches. All devices transmitting any traffic over the network are included in the discovery, ensuring that the data collected is complete regardless of when or for how long applications run.

How it's Used

Initial discovery. Cascade Discovery identifies all the CIs on your network and their dependencies to provide an accurate and up-to-date inventory for your CMDB. This inventory helps you prioritize implementation of scan-based discovery to complement the breadth of data supplied by Cascade with the depth that scanning provides where needed.

Scheduled CI and dependency updates. Cascade Discovery continuously collects data about CIs and their dependencies, including changes, and exports that data on a scheduled basis to BMC Atrium to ensure that your CMDB is always accurate.

Optimize scanning. The information that Cascade Discovery provides enables you to ensure that all systems that should be scanned are being scanned and that new scanning candidates are being appropriately selected based on specific needs. This enables you to get the detailed information scan-based discovery provides for all and only those systems where it is needed to minimize performance and credentialing challenges.

Optimize use of agents. The data center's dynamic nature can result in servers without appropriate coverage. Cascade Discovery enables you to automatically identify hosts in the data center and apply agents as appropriate to accurately assess those assets. Many critical assets also exist outside of the data center; Cascade Discovery identifies these assets, enabling you to apply the appropriate level of identification and management to the assets.

Deployment

The appliance can be deployed to discover CIs in a subset of your network or across the entire enterprise; in either case, you will have comprehensive CI inventory and dependency information within hours.

Key Features

Continuous, network-wide discovery. No gaps in time or network coverage.

Passive deployment. No credentials or agents needed.

Complete dependency mapping. Understand actual dependencies based on real-time usage and business service priority.

Fast deployment. Uses network flow from your existing routers and switches.

Riverbed Cascade

125 CambridgePark Drive
Cambridge, MA 02140
(617) 354-9292
www.riverbed.com/cascade

Riverbed Cascade EMEA Office

No 1, The Courtyard
Eastern Road
Bracknell
Berkshire RG12 2XB
UNITED KINGDOM
+ 44 1344 354910