APPLICATION MAPPING FOR MULTI-CLOUD ENVIRONMENTS: START HERE, START THERE, START ANYWHERE





map? Application maps provide clear visibility into which parts of the IT infrastructure support a business service, from applications in a traditional data center to the various components of your multi-cloud environment.

An application map makes it possible to document and monitor business services. It enables IT to assess the impact of potential changes, check service continuity configurations, and check other functions to ensure optimal business support.

But creating an application map is not a trivial exercise, especially in today's diverse multi-cloud environments. Application owners have to manage complexity. There are a lot of moving parts tied up with a modern, distributed business application: network infrastructure, databases, storage, servers, and software, not to mention applications and services delivered via public, private, and hybrid clouds.

It's common for the data center team to have an inventory of hardware and software assets. But that inventory doesn't reveal how these assets connect to one another and how they support an application.

When data for application mapping is collected and entered manually—which is often the case—by different IT functions—the risk of inaccurate data rises. Plus, it takes so long, by the time the data is collected and consolidated, it is likely out of date. This is especially true given the speed and pace of today's cloud deployments.

As a result, you end up with an inaccurate map, which leads to doubt of the quality of your application map. This has downstream affects as an application map is used as the basis for event correlation, impact assessment, <u>cost transparency</u>, and giving <u>business context to security</u> <u>vulnerabilities</u>.

Auto-discovery and Application Mapping

So manual application mapping is a dead end. What about auto-discovery tools?

These tools collect data on network infrastructure, servers, software, cloud services, and their relationships to build robust, tiered application maps. Auto-discovery also makes it feasible to easily update maps to reflect changes, which is especially critical in a dynamic multi-cloud environment and ensures timeliness and accuracy. Flexible reporting options accommodate a variety of documentation needs as well.

By fostering collaboration between application owners and configuration managers and extending across all your data center assets and cloud services, application maps advance IT-business alignment and help keep the environment responsive to business demands.

Auto-discovery is Not Enough

On their own, auto-discovery tools do not solve all the problems application owners face when trying to generate comprehensive, accurate maps. Where the mapping process starts can make a big difference in how accurate the map is.

Sometimes auto-discovery tools use a "top down" approach. That means the person creating the map has to specify the "top" entry-point for the application—often a URL that leads to a load-balanced service or to a web server. The choice of what to scan is *also* driven top-down. At each

stage, the tool decides what to scan based upon the dependencies it detected at the previous stage.

But there's not always just one entry-point. Application owners often think about where applications store data, which is akin to starting at the bottom. For off-the-shelf applications like Exchange, the natural place to start is with the Exchange software, rather than either the entry-points used by clients or the databases storing the data. Exchange servers are in the "middle," with entry-point servers and load balancers above and databases below.

Start-Anywhere Application Mapping

Start-anywhere application mapping is the only effective way to accurately map diverse and complex applications across physical and virtual infrastructure, including multi-cloud environments. The start-anywhere approach gives application owners a better experience by allowing them to start with what they know, rather than forcing them to work out what constitutes the "top."

<u>BMC Helix Discovery</u> always looks everywhere, so it sees systems and their dependencies from all directions. It removes the key limit of the top-down approach, which is that often a dependency between two things can only be seen from one end. In those cases, the top-down approach only sees the dependency if it's visible from the top.

Start-anywhere mapping can also start from multiple points simultaneously, which means you can cope much better with applications when parts or relationships are missing. Some dependencies will remain invisible to any tool, but with a rigid top-down approach, as soon as the tool meets something it can't discover, that's the end.

With the start anywhere approach, you look from the top, bottom, and middle—pinpointing and filling in the gaps much more quickly. This mitigates many of the challenges that standard tools face in today's complex, cloud-centric infrastructures.

How does it work? Users employ a modelling editor to create the map. They can annotate and trim a map—making it easier to read. For example, a user might turn off the feature that shows software components associated with software instances. <u>BMC Helix Discovery</u> can accomplish all this with ease in the most complex multi-cloud environments, and if the modeled environment changes, the map changes automatically.

Start-anywhere application mapping: just one of the ways that <u>BMC Helix Discovery</u> creates a holistic view of your application dependencies, from the data center to all your clouds and beyond.

Learn Why Discovery is Critical to Multi-Cloud Success