

---

# End Your CMDB Saga

**with Automated Discovery and Dependency Mapping**

---

An eBook from

**RESOLVE** 



# Executive Summary

The Configuration Management Database (CMDB) has been around for decades, creating heartache and misery for the IT people tasked with maintaining it. Originally implemented to help track and manage ever-changing IT assets, the CMDB is usually out of date the moment after its last update.

This eBook explores the checkered past of the CMDB, and how new technologies can help it overcome a bad reputation, ushering in an era where the CMDB finally delivers its long-awaited promise.


We'll explore:

- Why the CMDB is the cornerstone of successful IT operations
- The origins of the CMDB
- How the CMDB has evolved over the last 40 years
- Why getting your CMDB right is so important (and so difficult)
- A framework and action plan for CMDB success
- Best practices for getting your CMDB right and keeping it right
- Leveraging new technologies to achieve accurate, up-to-date data
- Laying the foundation for intelligent IT automation initiatives

We hope the following insights will help you finally end your CMDB saga — and provide an entertaining history along the way.

## RESOLVE POLL

Which best describes your organization's CMDB?



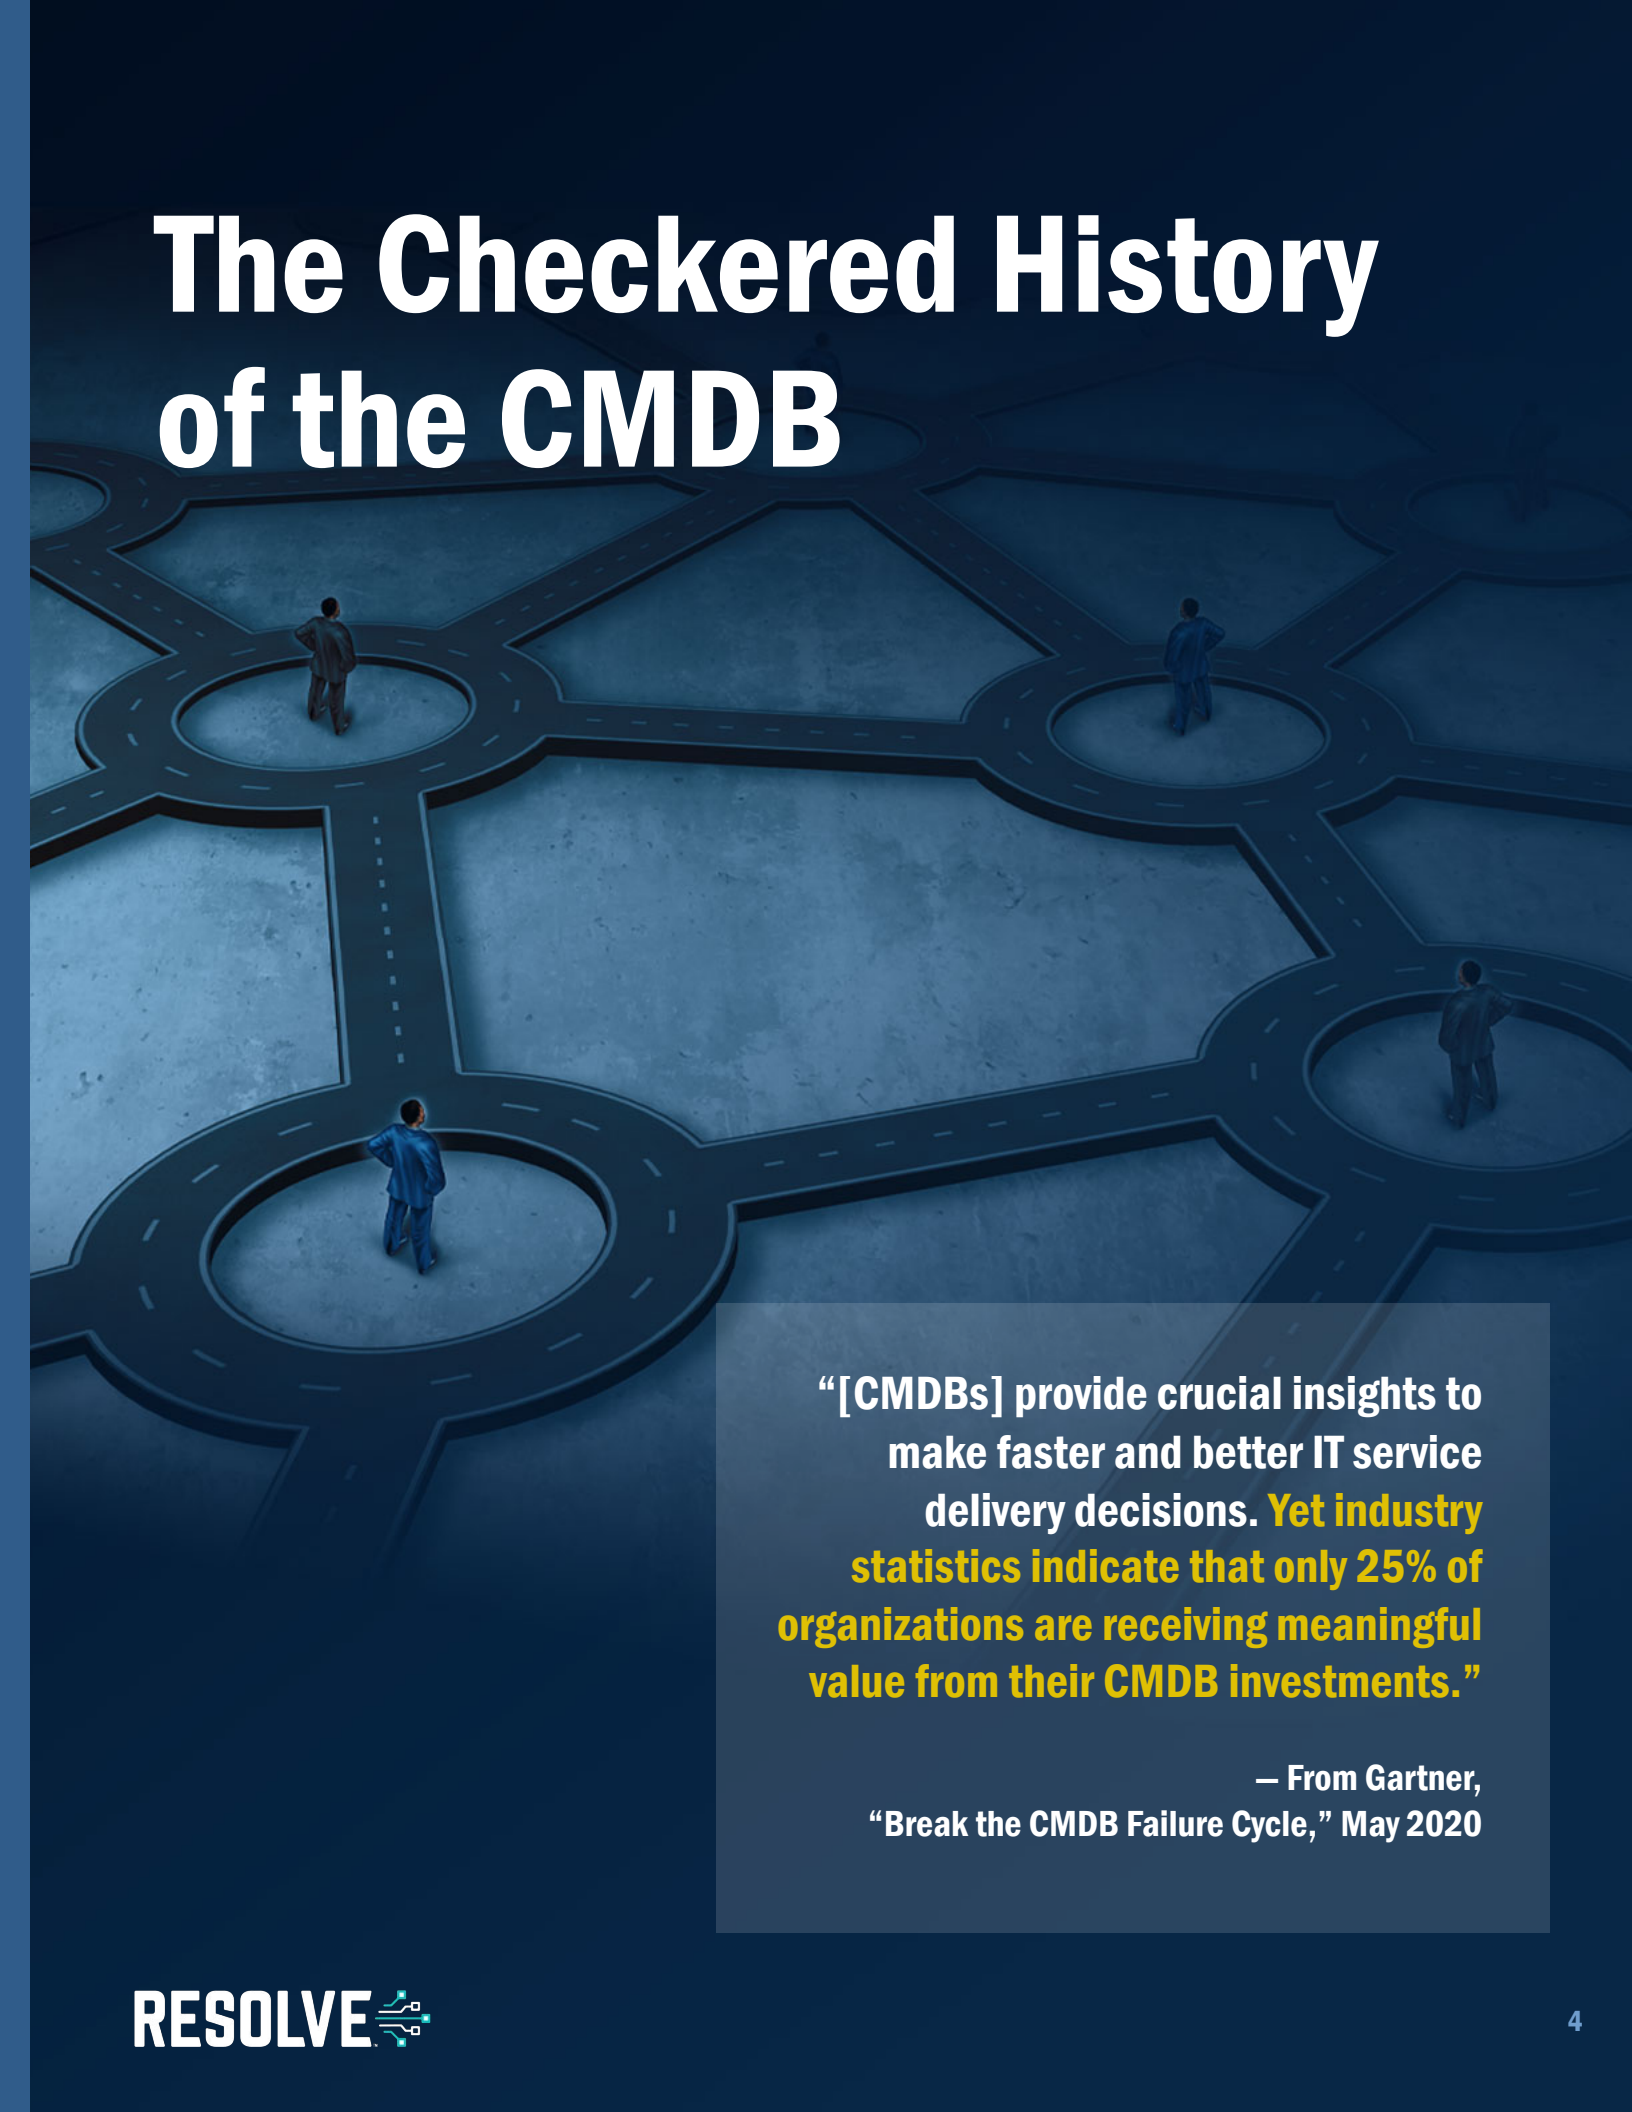
# CMDB: The Cornerstone for Successful IT Operations

A well-maintained, accurate CMDB is the cornerstone of successful IT operations, offering clarity about what's in your environment and how it all works together. This information is critical to maintaining efficient management, effective troubleshooting, mitigating change management risk, accurate planning, and more.

When done right, your CMDB should tell you:

- The complete inventory of available resources in your environment and what's running on them
- Resource locations and interconnectivity
- How each resource supports the business
- Resource cost and utilization
- Resource security
- How changes will impact critical applications, infrastructure , and budget
- How to quickly resolve problems when they occur (and find the root cause)

# The Checkered History of the CMDB



“[CMDBs] provide crucial insights to make faster and better IT service delivery decisions. **Yet industry statistics indicate that only 25% of organizations are receiving meaningful value from their CMDB investments.**”

— From Gartner,  
“Break the CMDB Failure Cycle,” May 2020



# 1980s: The Beginnings of the CMDB

Soon after the very first IT devices were deployed, IT professionals realized they needed a way to track them. They also needed a record of device repairs and updates, and a means to determine what else might break as a result of those changes. Thus, the Configuration Management Database (CMDB) was born – and immediately became outdated and despised by IT teams everywhere.

The concept of the CMDB first emerged in the 1980s as part of the IT Infrastructure Library (ITIL), a best practice framework produced by the UK government to help manage and develop controls for IT services. The ITIL principles and standards called for creating and maintaining a database to track and manage IT services. Conceptually, the CMDB is nearly as old as the ITIL standards themselves, and it is still considered a foundational element for IT Service Management (ITSM).

The CMDB was introduced as a process in ITIL V2 after IT pros realized that asset inventory wasn't cutting it, especially for change control. The concept didn't come about overnight, though. It emerged from the asset management process, which evolved into hybrid configuration management and then change control. These incremental improvements resulted in a true configuration management database.

Despite the need for an accurate CMDB, the practice suffered from inherent challenges. CMDBs were too complex, too cumbersome, and too difficult to maintain. In most cases, the CMDB only represented a snapshot in time; within days or weeks, it was out of date and worthless.

Recently however, the CMDB has been making a comeback from innovations in application and infrastructure discovery and dependency mapping (DDM) - more on that later.



# 1990s: The Era of Asset Management

In the 1990s, IT teams leveraged asset management tools to inventory physical hardware and software assets. At the time, physical assets were very expensive and, as they proliferated, organizations needed to track their investments.

Additionally, asset management was critical to validate software licensing and perform asset depreciation from an accounting perspective. Before true virtualization, logical partitions (LPARs) enabled physical machines to be split into multiple partitions, each hosting a separate instance of an operating system and software. LPARs were also tracked in the CMDB, and then once virtualization went mainstream, virtualized assets and the software running on them were recorded to support true ups and allocations. As complexity increased, asset management attempted to keep up.

As the CMDB evolved, IT pros started asking additional questions, like how assets were used and how they were connected to one another. In the case of a physical server, they wanted to know what software was running on the hardware, and how the company used the asset. And thus began the next phase of the CMDB evolution.




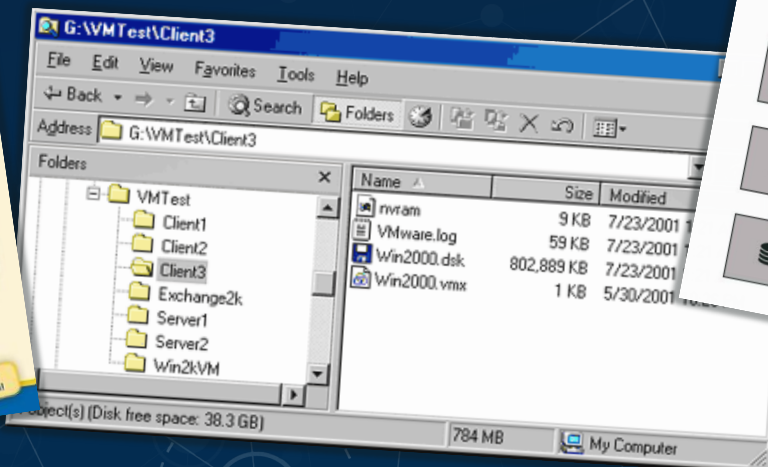

# 2000s: Configuration(ish) Management

In the 2000s, trends and developments in the CMDB morphed into configuration – taking stock of all the assets that existed and grouping them by relationships. As virtualization was embraced, there was an increasing need to understand the correlations between assets and gain visibility into what was happening upstream and downstream from each device and how those devices related to one another.

Teams initially leveraged a three-tier stack – comprised of a database sitting on an application server tied to a web server – to monitor critical business services. For example, an HR application would sit on top of an Oracle database tied to an application server and web server. If the web server failed, so did your HR app. This was the beginning of mapping and relationships in a configuration approach.

Around the mid-2000s the true CMDB finally emerged, bringing change control to the party. Change control offered additional insights to improve IT management and ensure there was a consistent record of every change made to each device to aid in troubleshooting, as well as compliance.

IT teams also started looking at infrastructure and apps through the lens of business services, providing another layer of visibility into how connected hardware and software came together to support business-critical applications and services.






# 2010s: From Peak of Inflated Expectations to Trough of Disillusionment

At the beginning of the 2010s, people had high hopes for the CMDB. Given its promise and potential to simplify IT operations, there was significant time and effort put into building out configuration databases in most enterprise organizations.

Unfortunately, experts soon realized that it was almost impossible to track everything in the CMDB because there was simply too much to input manually. Adding to that pain, the CMDB was only able to capture upstream-downstream relationships while left-right relationships did not translate very well.

By the mid-2010s, the CMDB had lost its mojo. It was too cumbersome and never accurate, and its reputation reached an all-time low.





# 2020s: Technology Developments Usher in a CMDB Renaissance

Recent technology developments have finally made automated discovery and dependency mapping a reality, enabling near-real-time updates to CMDBs that deliver the long-awaited accuracy required to fulfill their potential.

By delivering full-stack visibility into dynamic, hybrid enterprise IT environments, IT teams can finally get the visibility they need to more efficiently and effectively manage infrastructure and applications.

Today's discovery technologies automatically identify and inventory all of the compute, network, and storage entities across data centers and multi-cloud environments. Dependency mapping then identifies the upstream-downstream flows to effectively map the relationships between assets and applications and form groups that make up critical business services.

All of this data is pushed to the CMDB in real time – without any human interaction required. What's more, automation enables updates to the CMDB's asset class structure while auto-creating changes and change requests.

Together, these advances mean that the CMDB can finally become the true cornerstone of IT operations, establishing a strong foundation for more advanced automation and intelligent IT operations.



# Why Is It So Difficult to Get Your CMDB Right?



“At least three vendors that we interviewed used the term ‘**death spiral**’ to describe CMDBs that didn’t manage data quality well.”

— From Forrester, “Rethink Your CMDB”  
September 2020

# CMDB Challenges

Anyone who has attempted to build or maintain a CMDB can attest to the many challenges IT pros face when it comes to taming the CMDB beast. Many organizations have taken multiple runs at getting their CMDB right.

Here are some of the most common obstacles we've encountered on the journey to CMDB success. If you understand the challenges from the outset, you can develop a game plan to overcome them.

## Traditional Challenges:

- Teams never trust the data because it is usually out of date, incomplete, or simply inaccurate.
- Manual updates just can't keep up in today's dynamic IT environments, especially those with multiple deployments per day.
- Mapping dependencies by hand is exceedingly time consuming and riddled with human error... and it's impossible in larger environments.
- Highly virtualized environments with containers and serverless architecture are difficult to track. It's not about physical devices anymore.
- Hybrid, multi-cloud, multi-domain environments are too complex to track manually.
- IoT means millions of devices (of many types) are now part of your IT environment.
- A lack of executive sponsorship is a common downfall of many CMDB initiatives.

## New Challenges Introduced by COVID-19:

- Who knows what's being pushed out and by whom?
- IT governance and controls are more challenging than ever before.
- Data gaps are more glaring than ever before.
- New devices need to be tracked and managed as workforces are increasingly remote and distributed.
- Everything is a web application.



# Framework & Action Plan for Success

“Many CMDBs have failed by having too much detailed data forced into them **without a clear business purpose or maintenance plan.**”

— From Forrester, “Rethink Your CMDB”  
September 2020



# Actionable Steps to Getting Your CMDB Right

Ready to finally get your CMDB right? First and foremost, we recommend that you develop a solid strategy before diving in. Determine what data you want to start with, how you plan to use it, and how it will be maintained over the lifetime of your CMDB. The data will quickly become useless if there's not a plan for it to be actioned and continually updated.

Here's a high-level blueprint to help you chart your course:



## IDENTIFY YOUR EXECUTIVE SPONSOR & DRIVE ALIGNMENT

Executive support and sponsorship is critical to the success of your CMDB initiative, as well as alignment across your stakeholders.



## DOCUMENT YOUR GOALS & CMDB VISION

Outline the long-term vision for your CMDB and how you plan to leverage the data to meet specific business goals.



## DETERMINE YOUR DATA STRATEGY

Identify the data you need to capture to support your goals, and which applications and services are most important. Set up the classes in the CMDB and ensure each class has an owner and process attached to it.



## INTEGRATE TOOLS TO AUTOMATE DISCOVERY & INFRASTRUCTURE, APPLICATION DEPENDENCY, & SERVICE MAPPING

Automate discovery and dependency mapping to capture and push data to your CMDB to ensure it is always accurate and up-to-date.



## KEEP YOUR EYE ON THE PRIZE

Ensure you stick with your process as tools are deployed and integrated; then validate throughout the entire lifecycle.

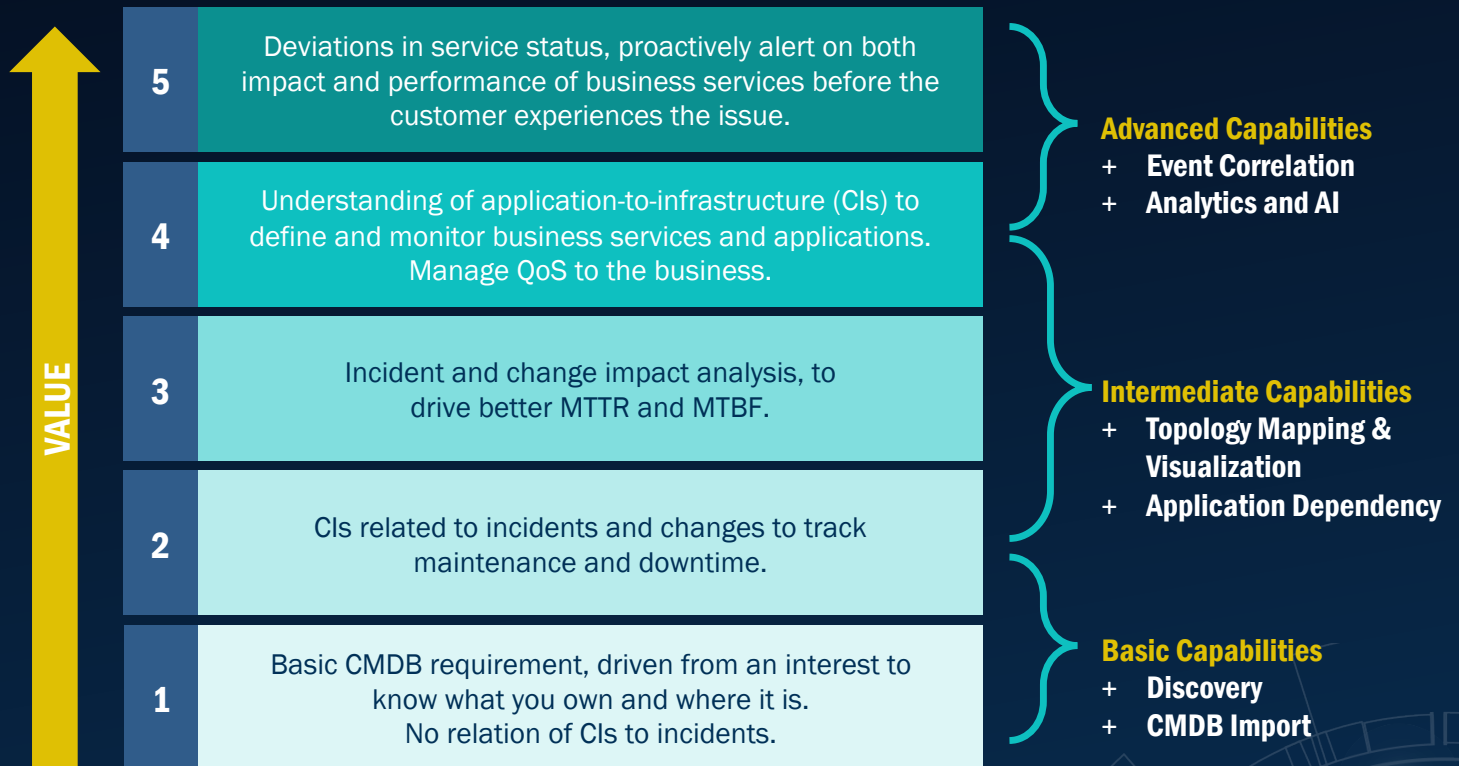


## ONGOING IMPROVEMENT & EVALUATING KPIs

As you define, measure, and begin to see results, you'll continue to improve and expand your process. The greater the detail, the greater the value of your CMDB.


# CMDB Maturity Model

How mature is your CMDB? This model will help you determine where you are today and how to take a phased approach to incremental improvements. As you work your way up the maturity model, the items on the right of the model will help you achieve the next milestone.



## RESOLVE POLL

How confident are you in your CMDB's accuracy?

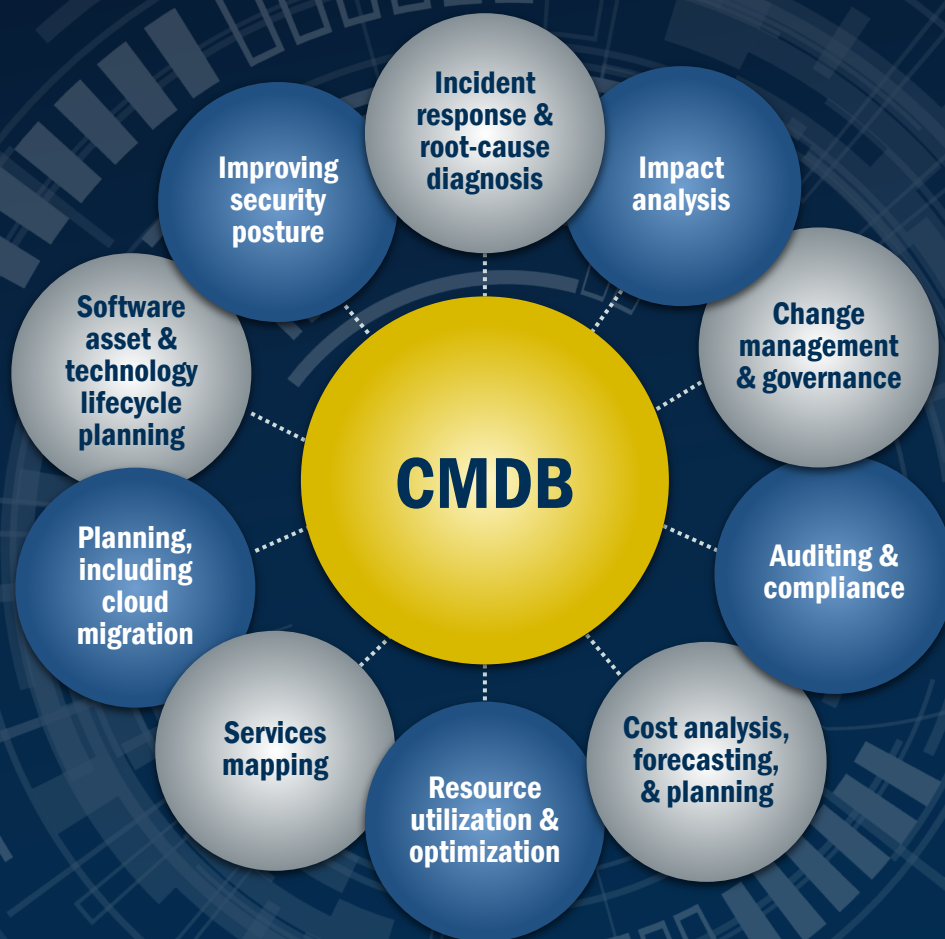




# Making Your CMDB the Center of Your Universe

As your CMDB evolves, it has the potential to become the center of your IT universe, fueling everything from change management and governance to underpinning the planning process for complex cloud migrations.

When teams have confidence in the CMDB's accuracy, they can use this data for both strategic decision making, as well as streamlining day-to-day tactical operations.



# New Technologies Can Help You Succeed

“Until now, I&O pros couldn't gain an accurate understanding of their enterprise application ecosystem. Attempts to fix this shortcoming failed due to the unscalable manual processes they required — the best-known failure being the rise and fall of CMDB... [today] successful **AIDM provides I&O teams with a big-picture topology of how applications and infrastructure connect**, making operational and business processes more efficient.”

— From Forrester, “It's Go Time For Application and Infrastructure Dependency Mapping (AIDM)”  
February 2020



# Automate Discovery & Dependency Mapping

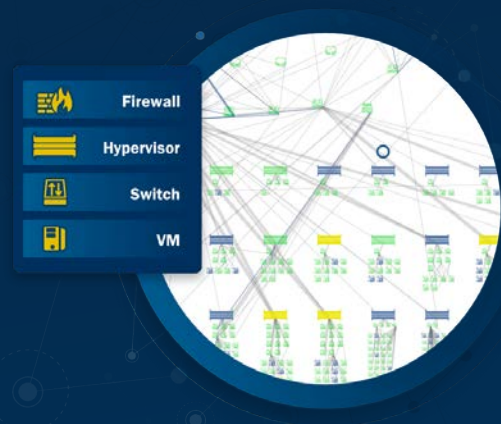
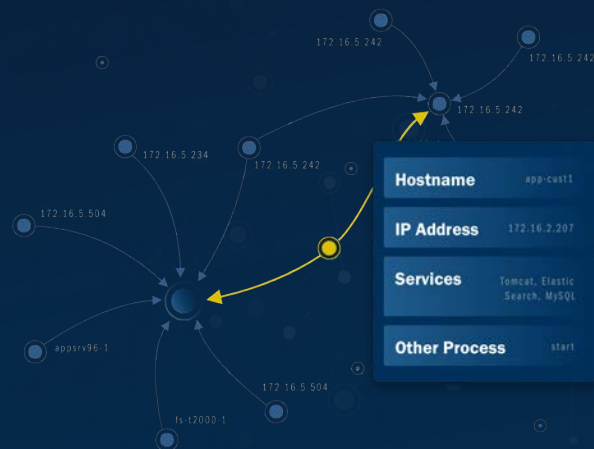

Automated discovery and dependency mapping can finally deliver the deep visibility you need into complex, dynamic IT infrastructure and ensure that your CMDB is always accurate. Automating these functions helps overtaxed IT teams by simplifying management, streamlining operations, and offering a unified view into multi-domain environments.

Resolve provides agentless auto-discovery and dependency mapping (DDM) as part of a broader suite of advanced IT operations capabilities. Resolve's DDM quickly discovers compute, network, and storage entities across dynamic, hybrid IT environments without deploying any agents.

Lightweight data collectors intelligently scan subnets or user-input boundaries to identify CIs and determine a variety of details about each entity, including vendor make, model, platform, configuration, IP address, and more. A real-time inventory is produced and can be filtered and viewed in a variety of ways.

Additionally, Resolve identifies and tracks dynamic, multi-layer dependencies between infrastructure components, applications and supporting infrastructure, and services. The resulting topology maps make it easy to visualize the relationships between CIs, applications, and services, providing unprecedented visibility for troubleshooting, change requests, and proactive maintenance. You can even overlay fault, performance alerts, and tickets right on top of the topology maps to identify root cause and improve MTTR.

Finally – and most importantly for the subject at hand – Resolve provides bidirectional connectors to your CMDB that enable auto-discovery and dependency mapping data to be automatically pushed to the CMDB in near real time. This ensures you always have an accurate, up-to-date record of everything in your environment and how it is all connected.





# How Will You Benefit?

While getting your CMDB right is a journey, it's worth the effort. In addition to offering many immediate benefits, it also sets the stage to advance digital transformation efforts and to accelerate your automation initiatives.

Here's a quick summary of the benefits you can expect from automating discovery and dependency mapping:



## FULL OBSERVABILITY

Auto-discover and map all of the network, compute, and storage CIs across hybrid environments



## IMPROVE MTTR

Accelerate incident response with deep visibility and root-cause analysis



## SAVE TIME

Eliminate countless hours of manual work and human error



## INCREASE PERFORMANCE

Improve application and infrastructure performance and uptime



## UNIFY YOUR VIEW

Access all your CI data in one place through a single pane of glass



## DERISK CHANGE MANAGEMENT

Make changes confidently and safely with full context



## ENSURE ACCURACY

Rest assured that your CMDB is always populated with the latest discovery and dependency details



## REDUCE COSTS

Eliminate unnecessary or underutilized CIs to save money



## GAUGE IMPACT

Quickly quantify the business impact of outages to ensure stakeholders are informed and issues are correctly prioritized



## IMPROVE COMPLIANCE

Identify unknown CIs that pose a risk and auto-generate reports for your auditors that offer the full visibility they require

## ABOUT RESOLVE

Resolve Systems helps enterprise technology teams worldwide achieve agile, autonomous operations with an industry-leading intelligent IT automation platform. With more than a decade of automation expertise, Resolve's solutions are purpose-built to address challenges posed by increasing IT complexity. Resolve enables organizations to maximize operational efficiency, overcome labor shortages, reduce costs, quickly troubleshoot and fix problems, and accelerate service delivery. Resolve is majority-owned by Insight Partners, a leading global venture capital and private equity firm investing in high-growth technology and software companies.

[Visit our website to learn more ›](#)

## Award-Winning Intelligent IT Automation Technology



Ready to see what DDM can do for you? [Request a demo.](#)