INTRODUCTION TO IT DISCOVERY

Today's IT organizations continue to grapple with the twin challenges of evolving to keep up with changing customer and technology needs while maintaining stability on existing services in the face of disruptions and risks.

For this reason, visibility into IT components. including their interconnections and dependencies, has become a key differentiator in successful service management practices. Without visibility, there is a real risk of business loss and dissatisfaction in the services you deliver to customers.

While IT discovery may seem like a piece of cake for smaller organizations, larger enterprises are finding this to be a challenge. This has been exacerbated by the following drivers:

- The sheer size of IT components that require integration, such as virtual machines (VMs), containers, APIs, and/or IoT devices.
- The increasing IT sprawl which has been impacted by shadow IT, BYOD, and remote working, especially in the pandemic era.
- The current tendency to adopt hybrid cloud architectures, leading to complexity as it becomes difficult to have a clear picture of what is deployed and how it fits into the puzzle.

What is IT discovery?

IT discovery is the process of identifying and collecting data on existing IT components within a network.
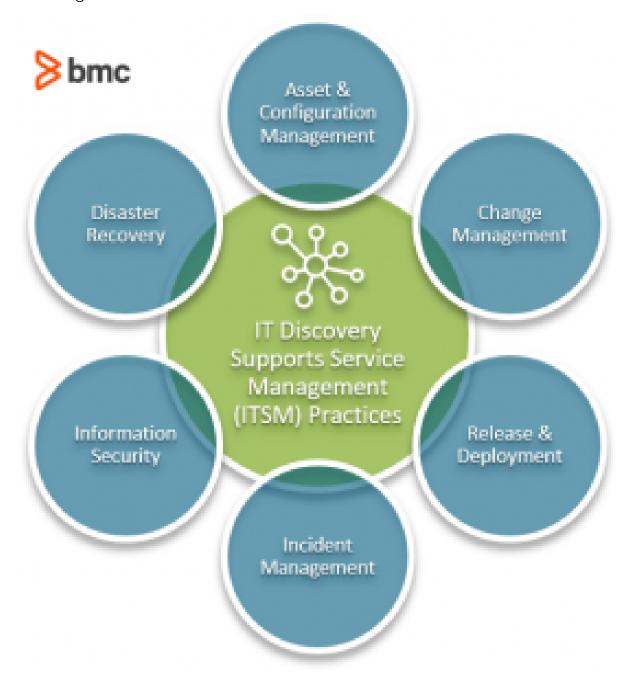
This involves carrying out a sweep across the network using IT tools to find all the IT assets within the specified environment including hardware, software, virtual instances, cloud components, and smart devices, among others.

The kind of information that can be gathered during a discovery includes:

- Device name, device type, network address
- Device configuration: hardware/software/firmware versions
- Device status, capacity, and performance information

Why do we need IT discovery?

Discovery allows us to take a snapshot of IT components or environment—their current state. This facilitates capturing of baseline information that can be used to support other ITSM practices, including:

Asset & Configuration management

Discovery is an important tool during audits of IT assets. You might unearth discrepancies between asset records held by Finance and those maintained by IT, especially where the status information is inconsistent e.g. devices that are not in use, lost, or out of service.

Discovery can also unearth devices that are using unlicensed software or incorrectly configured, which can cost the organization from a cost or effectiveness perspective.

(Learn more about [asset management](#) & [configuration management](#).)

Change management

Technical teams reviewing changes to determine potential impact on other services and components can leverage insights provided by IT discovery data. Such insights can:

- Prevent unplanned incidents
- Identify dependencies that will require involvement of other stakeholders in making the change successful

(Know [types & levels of change management](#).)

Release & deployment

Discovery can aid in planning of releases, as visibility into whether dependencies such as libraries and OS versions can be determined in advance and validated before installation of new or updated software components in the IT environment.

Discovery can also be employed during post-verification checks after deployment activities.

(Compare [deployment to release](#).)

Incident management

Investigations into incident causes can be aided by discovery information, as information on service and component configuration, status, and dependency can be very valuable for technical teams involved in resolution efforts.

(See how [incident management](#) should work.)

Information security

Discovery can be useful in [mobile device management](#) particularly in identifying and managing BYOD where data breach risks abound from personal owned devices accessing business and customer information.

Discovery is also a key capability used for identifying access points when enumerating a network for purposes of [penetration testing](#).

(Explore [information security](#), which starts with asset discovery.)

Disaster Recovery

Following a major disruption that might require an organization to restore its services in a different IT environment, information from discovery can prove an invaluable resource in rebuilding IT services and components from scratch.

(Plan for disaster with [these best practices](#).)

How does IT Discovery work?

At the most basic level, discovery of active IT devices on a network can take place through Address Resolution Protocol ([ARP](#)) commands. The ARP cache contains one or more tables that are used to store IP addresses and their resolved Ethernet or Token Ring physical addresses. ARP works for IPv4, but has since been replaced by Neighbor Discovery ([ND](#)) for IPv6 devices, which works in a similar fashion.

Beyond ARP, [SNMP](#) (Simple Network Management Protocol) is the next level of discovery. Devices that are configured for SNMP can respond to messages called protocol data units (PDUs) from devices within the same network segment. The discovery tool will send an SNMP GET message requesting for information from other devices that are SNMP enabled. The information is stored in Management Information Bases ([MIBs](#)) that contain configuration and status information.

Other approaches for IT discovery include:

- **Common Information Model ([CIM](#)).** This open standard provides a common definition of management information for systems, networks, applications, and services, and allows for vendor extensions. It is a modeling schema that describes managed system, hardware, and software objects.
- **Windows Management Infrastructure ([MI](#)).** Based on CIM, MI is an evolution of Windows Management Instrumentation ([WMI](#)), a set of specifications from Microsoft for consolidating the management of devices and applications in a network from Windows computing systems.
- **Cloud discovery**. This approach involves use of APIs to enumerate cloud instances and gather information on cloud configuration.

Specialized IT discovery tools are able to process the information in MIBs and create topologies of network architecture as well as display status of various network elements. These tools can be triggered via manual action or configured to automatically search and update asset information. A common distinction in how these tools work is the agent vs agentless approach:

- **Agent-based approach.** Here a small piece of software is installed in each component within the network to collect and send back information to the central discovery tool on availability and performance metrics. Agent-based discovery provides greater visibility as the level of detail better supports monitoring and troubleshooting efforts. However, it comes with the overhead of installation, updating and management, which can be automated but still requires significant administration effort.
- **Agentless approach.** Unlike the former, here no installation is done on components to be discovered. The discovery tool uses the previously mentioned protocols to discover status and performance of IT assets, with reduced administrative overhead. The amount of detailed information from the discovery process is obviously less than agent-based discovery.

IT discovery isn't about the tooling

IT discovery shouldn't be treated as just an exercise to turn on a tool, let it gather data, and that's it. It is critical that you define a discovery process that ensures collected information is both:

- Analyzed to make sense of it.
- Availed to the right people to act on it.

Value can only be generated from IT discovery if it supports the ITSM practices mentioned earlier, leading to better managed services, cost effectiveness, and overall business success.

Related reading

- [BMC Service Management Blog](#)
- [Automatically Discover Assets & Their Relationships](#)
- [IT Security Policy: Key Components & Best Practices for Every Business](#)
- [Enterprise Networking Explained: Types, Concepts & Trends](#)
- [Data Center Migration: Creating a DC Inventory](#)