


precisely

Guide to Mainframe and IBM i Configuration Items in ServiceNow®




Digital Transformation and IT Operations Management

Digital transformation has spread to virtually every industry and country around the world, with spending on related technologies and services expected to top \$2.3 trillion globally by 2023 (IDC). As a result, IT infrastructure has exploded in breadth and complexity, and organizations are highly dependent on that infrastructure to keep business up and running.

The criticality of IT infrastructure to the business, the need to ensure its security, and the cost associated with deploying and maintaining it, have driven many enterprises and government agencies to invest in IT operations management (ITOM) solutions. These solutions enable organizations to get better visibility into their infrastructure and services in order to maximize operational agility and ensure high service availability, security and compliance for the business.

ServiceNow is an industry leader in IT Workflow solutions, including ITOM, with a suite of applications that include discovery, event management, orchestration, operational intelligence, service mapping and cloud management. At the heart of all these applications is the ServiceNow Configuration Management Database (CMDB).



Configuration Management and CMDB

Think about all the components – servers, routers, switches, desktops, and many more – that make up the far-reaching and complex IT infrastructures that have rapidly evolved as a result of digital transformation. Each one serves a purpose in supporting your business and requires some degree of visibility and management. Not only are these components important, but they, in turn, have subcomponents – and these can also be important to see and manage.

This is where configuration management comes in. ITIL defines configuration management as:

The process responsible for ensuring that assets required to deliver services are properly controlled, and that accurate and reliable information about those assets is available when and where it is needed. This information includes details of how the assets have been configured and the relationships between them.

A configuration management database (CMDB) is the database built specifically to store the information you collect about your hardware and software assets, including the relationships between them. The configuration item (CI) is the building block of the CMDB. A CI can be any application, infrastructure, or service component that you are managing. CIs are grouped by classes, and each CI has attributes to describe the component. For example, a server CI would likely have CPU and memory as attributes.



What to Include in Your CMDB ... and How to Do It

ServiceNow recommends that you start simple when building a CMDB, and then expand as your organization and configuration management processes mature. Starting with a strategic initiative and use case can help guide which CIs, and which attributes, are important to include in the CMDB first. For example, if improving the performance and availability of a customer online payment application is a strategic initiative, you might focus your initial CMDB population with all the components required to fulfill that application end-to-end.

Once you've designed your configuration data model, you need to populate your CMDB – and keep it up to date and accurate. If you don't have a plan for keeping the CMDB current, all the time and effort spent setting it up will be wasted.

Manually entering information into the CMDB is extremely timeconsuming and error-prone, and keeping it up to date as CIs change is nearly impossible to sustain over the long term.

ServiceNow addresses this challenge with its ServiceNow® Discovery application, which automates the discovery of all your physical and virtual devices, such as laptops, desktops, servers (physical and virtual), switches, routers, storage and applications. It also discovers the dependent relationships between them. Discovery runs on an on-demand or scheduled basis to help ensure that the completeness, accuracy and currency of the CI data in your CMDB.



Don't Leave Out IBM i and Mainframe Systems

Discovery greatly accelerates the time to value of ServiceNow customers building a CMDB and keeps it up to date so that you can be confident in the CI data underpinning all your ServiceNow applications across the enterprise. In other words, the CMDB can be considered a shared, single version of truth.

But there's a missing piece of the puzzle. ServiceNow does not natively discover CIs from IBM i systems and IBM mainframes.

However, banks, healthcare, financial services, retail, government and many other organizations continue to rely upon IBM i systems and IBM mainframes to run mission-critical applications and databases. Only by incorporating data from these platforms into the ServiceNow CMDB can these enterprises ensure the integrity and reap the full value from their ITOM initiatives.

If you have an IBM i or mainframe in your organization, it is important you can identify and understand the CIs to include in your CMDB. The following sections provide a list of the configuration item classes you will want to capture.



IBM i Configuration Item Classes

Approximately 16 CI classes, consisting of more than 156 attributes, are commonly discovered in the IBM i environment. Each class includes one or more attributes that reference other CI classes, which are important in documenting the relationships and dependencies among configuration items.

A sampling of these classes is organized into three categories below: the IBM i Frame, LPAR and Object Discoveries.

IBM i Frame

Frame: The containment unit that physically houses nodes, switch hardware, and other control/supporting hardware. For the frame, you will want to discover basic attributes such as processor group/feature/feature code, interactive feature and memory.

Processors: You will want to also capture attributes specific to the frame processors, including location, card position, part number and more.

LPAR

Logical Partition (LPAR): A virtual server with a subset of the physical server's processors, memory, and I/O adapter with its own operating system instance and applications. For the LPAR itself, you will want to capture operating system name and version, as well as primary IP address, and memory.

Other CI Classes you will want to capture for each LPAR include:

Active jobs: Job is the basic unit of work on an IBM i system. There are several types of active jobs, including autostart, batch, interactive, prestart, subsystem and system jobs.

Auxiliary Storage Pools (ASPs): Auxiliary storage is all addressable storage (memory) other than main storage. Auxiliary storage pool (ASP) is a group of disk units defined from the auxiliary storage devices. An ASP provides a way of organizing data to limit the impact of storage device failures and to reduce recovery time.

Disk: Storage device that includes one or more flat, circular plates with magnetic or optical surfaces on which information is stored.

Installed software: For the software installed on the LPAR, you will want attributes including usage count, limit and peak usage.

Network connections and interface: Information to collect on the network include the connection type/connection status, interface type/name, as well as IP address (local/remote) and network address.

Subsystems: To efficiently use system resources, different types of jobs require different processing instructions and system resources. To meet this need, the operating system creates unique operating environments called subsystems. You will want to capture the subsystem library/name, as well as a description and maximum number of jobs allowed.

System values: Attributes including description, shipped value, current value, system value.

Object Discoveries

An object is a named unit that exists (occupies space) in storage, and on which operations are performed by the operating systems. IBM i objects provide the means through which all data processing information is stored and processed by the IBM i operating system. Through objects you can find, maintain and process your data on the system.

While there are unique attributes you will want to discover for each CI class, there are a number of common ones to maintain in your CMDB, including: Name, Description, ASP Number, Library, Object Size, Creation Date/Time, Creator User Profile, Owner User Profile, and Last Change Date/Time.

Files: Collection of related data that is stored and retrieved by an assigned name. Can be a database file, a device file, or a save file. The system-recognized identifier for the object type is *FILE.

Job queues: An object that contains a list of batch jobs waiting to be processed by the system. The system-recognized identifier for the object type is *JOBQ. Attributes include creation date/time, creator and owner user profile, job queue library/name/status.

Libraries: A library is a system object that serves as a directory to other objects. A library groups related objects, and allows users to find objects by name.

Output queues: An object that contains a list of spooled files to be written to an output device, such as a printer. The system-recognized identifier for the object type is *OUTQ.

Programs: A sequence of instructions that a computer can interpret and run without a user's intervention.



Mainframe Configuration Item Classes

An LPAR (logical partition) is a virtual server with a subset of the physical server's processors, memory, and I/O adapter with its own operating system instance and applications. For the LPAR itself, you will want to capture a variety of attributes including the LPAR name and ID; IPL (initial program load) date, time, device and volume; OS Release; and memory.

Other CI Classes you will want to capture include:

DASD storage: Direct access storage device – also referred to as disk. Attributes include tracks, free tracks, free extents, largest free extent, percent used, data set control block (DSCB), and the reference LPAR.

Jobs – z/OS jobs and active jobs: A mainframe job is the unit of work that a mainframe operator, or a job scheduler, gives to the operating system. Attributes include Job name and type, service class, user ID, workload name and service class. For active jobs, you will want to capture the position, step name, proc step, performance group number, program name, priority, and current storage usage.

Network CIs – network adapter, network connection, network routing; TCP/IP stacks: Information to collect about the network include the IP family, TCP/IP names, gateway, interface, flags, destination, bytes sent, protocol, connection time, state, and more.

Software – registered and active: While IBM software is most commonly run on the IBM Z mainframe systems, it is important to discover, and load into the CMDB, registered and active software from IBM as well as other vendors.

Storage Management Subsystem (SMS): Automates the use of storage for data sets.

Sysplex: A sysplex is a tightly-coupled cluster of distinct, independent instances (aka images) of the z/OS operating system that cooperate to process work.

Db2

IBM Db2® for z/OS® is a relational database management system (RDBMS) that runs on the mainframe to collect, store, organize and analyze data. Db2 provides users with the ability perform transactional, analytical and warehousing functions with the data contained in the database.

Related CI Classes include the following:

Database: Db2® databases are a set of Db2 structures that include a collection of tables, their associated indexes, and the table spaces in which they reside.

Data sharing groups: A collection of Db2 subsystems that access shared Db2 data.

Distributed Data Facility (DDF): Enables data access at Db2 servers by client applications running in an environment that supports distributed relational database architecture (DRDA).

Subsystem: Db2 operates as a formal subsystem of z/OS, a secondary system that is a distinct instance of a relational DBMS. Its software controls the creation, organization, and modification of a database and the access to the data that the database stores.

Tables: Tables are logical structures that Db2 maintains. They are made up of columns and rows. At the intersection of every column and row is a specific data item, called a value.

Table spaces: A Db2 table space is a set of volumes on disks that hold the data sets in which tables are actually stored. Every table is stored in table space.

CICS and IMS

According to IBM, most of the world's banking, insurance online trading, and order entry transactions are handled by the IBM CICS® (Customer Information Control System) and/or IBM IMS (Information Management System). With so much of today's critical digital infrastructure reliant on these two systems, it is essential that they be included in the ServiceNow CMDB.

CICS

CICS CI Classes to capture, include:

Groups and lists: Information on the CICS system definition (CSD) file is organized into groups and lists. The group collects related resources together on the CSD file. Every resource that you define must belong to a group; you cannot define a resource without also naming its group. A list contains the names of groups that CICS installs.

Regions: A CICS® region is a collection of resources that are controlled by CICS as a unit. The resources in a region include programs, Basic Mapping Support (BMS) map sets, transactions, terminals, files, transient data queues, temporary storage queues, and journals. CICS regions implement transaction processing services, which enables application programs to concentrate on business logic and not on how the logic is implemented.

CICSplex: A CICSplex is a cluster of CICS regions that can communicate with each other and cooperate to handle inbound work requests.

Transactions and programs: CICS applications are traditionally run by submitting a transaction request. A transaction is a unit of processing initiated by a single request. CICS transactions are used to perform multiple operations in the CICS region. Execution of the transaction consists of running one or more application programs that implement the required function.

IMS

Databases: The IMS database stores data using a hierarchical model, implemented using blocks of data (segments), which, in turn, can contain several pieces of data, called fields. IMS databases can be in two general classes: full-function and Fast Path.

Database areas: Subset of a data entry database (DEDB) containing one or more data sets.

Regions: The IMS control region provides the central point of control for an IMS subsystem, and provides all logging, restart and recovery functions for the IMS subsystems.

Subsystems: IMS runs as a z/OS subsystem.

Transactions: A transaction is a specific set of input data that triggers the execution of a specific business application program (a process or job). The message that triggers the application program, and the return of any results, is considered one transaction..

MQ

As is the case with CICS and IMS, IBM MQ drives business in today's digital world. According to IBM, 70 percent of the Fortune Global 500 rely on the IBM MQ messaging and queuing middleware to connect applications, systems, services and files across multiple platforms.

With IBM MQ, programs communicate by sending each other data in messages, which are placed in queues, so that programs can run independently of each other. MQ Channels are objects that provide a communication path from one queue manager to another.

When populating your ServiceNow CMDB, you will want to discover all three CI Classes, and their associated attributes:

- MQ Managers
- MQ Queues
- MQ Channels



Discover IBM i and Mainframe Environments with Ironstream and ServiceNow Discovery

As we established above, it is essential to include IBM i and mainframe systems in your ServiceNow CMDB – yet, ServiceNow doesn't natively discover the extensive number of CI classes and attributes that exist in these environments.

Manual and homegrown efforts are extremely complex, time-consuming and unsustainable. Organizations that have gone down this path report spending months of their IT team's time to populate the CMDB with their mainframe environment, only to have it be out of date once they were finally complete.

Precisely's Ironstream for ServiceNow software is a Now® Certified solution created to overcome these challenges by including these traditional IBM systems in the ServiceNow platform. Ironstream seamlessly integrates with ServiceNow Discovery to identify and build an inventory of IBM Z and IBM i CIs – along with their interrelationships, populate this information in the ServiceNow CMDB, and then automatically record changes and update ServiceNow. With Ironstream, organizations can achieve a complete, accurate, and up-to-date CMDB, without manual intervention.

To learn more about Ironstream for ServiceNow, visit: <https://precisely.com/integrate>





About Precisely

Precisely is the global leader in data integrity, providing accuracy and consistency in data for 12,000 customers in more than 100 countries, including 90 percent of the Fortune 100. Precisely's data integration, data quality, location intelligence, and data enrichment products power better business decisions to create better outcomes. Learn more at [precisely.com](https://www.precisely.com).

www.precisely.com

Copyright ©2020 Precisely. All rights reserved worldwide. All other company and product names used herein may be the trademarks of their respective companies.