WHITE PAPER

Visibility to Multi-Cloud

Measurable outcomes from CMDB

servicenow

Table of Contents

Journey to Cloud starts with ServiceNow ITOM

Cloud and DevOps are independent but mutually reinforcing strategies for delivering business value through IT. A hybrid or multi-cloud workload deployment offers the advantage of high resiliency, combined with the agility to adapt quickly to changing digital business requirements. Multi-cloud is quickly becoming the de-facto deployment standard as organizations of all types leverage an ever-increasing variety of cloud computing services.

Key factors that influence multi-cloud deployment strategies are:

- “Best-of-breed” service offering
- Cost to IT Ops and enterprise licensing
- Choice of technology stack for IaaS/PaaS/FaaS services
- M&A and data sovereignty, which plays a vital role on the decision-making process

Visibility into multi-cloud/hybrid cloud

For many customers, migration to cloud is a transformational journey. Visibility to multi-cloud/hybrid cloud deployment data with on-premise infrastructure and application data is critical to solve the real-world challenges from IT Operations. The ServiceNow platform has evolved as the de-facto standard for managing IT Ops business use cases, and the ServiceNow CMDB has emerged as the data platform for CloudOps. The CMDB is the single source of truth for any data requirements related to ITOps or CloudOps use cases. The Discovery product from ServiceNow offers near real-time visibility to multi-cloud IaaS/PaaS/FaaS services.

> “
> Multi-cloud is quickly becoming the de-facto deployment standard as organizations of all types leverage an ever-increasing variety of cloud computing services.

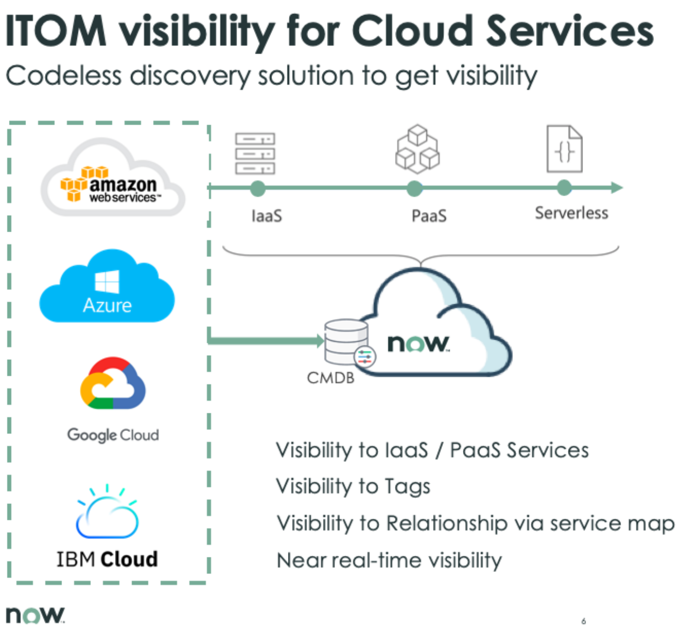

ITOM visibility for Cloud Services

Codeless discovery solution to get visibility

Visibility to IaaS / PaaS Services

Visibility to Tags

Visibility to Relationship via service map

Near real-time visibility

The breadth and depth of cloud discovery capabilities help to collect deployed cloud services data in the CMDB in near real-time for solving outcomes such as:

- **Service Management:**
  Incident/Change and Problem Management for workloads in cloud
- **Multi-Cloud Management:**
  Day 1 and Day 2 operations for cloud workloads
- **Software Asset Management:**
  Visibility into licensable software deployed in cloud and optimized software spend, based on vendor licensing agreements
- **SecOps:**
  Data exposure highlights sustained risk from poor information-protection practices on cloud workloads.[1] Incident response requires near real-time visibility to automate the prevention of security threats. Vulnerability response requires visibility into installed software and its vulnerability in cloud workloads.
- **Cloud cost optimization:**
  Reduce software spend by optimizing cloud workloads. It's difficult to manage IT spend across a hybrid IT landscape. Visibility into cloud workloads and reconciling the data with the cloud billing data enables transparency and provides Showback on cloud consumption.

> “
> Visibility into cloud workloads can be accelerated by service account-based [API-BASED] Discovery.

Customer journey to Cloud Discovery

Visibility into cloud workloads can be accelerated by **service account-based [API-BASED] Discovery**. Customers can get visibility into IaaS, PaaS, and FaaS services simply by adding their service accounts to the ServiceNow platform.

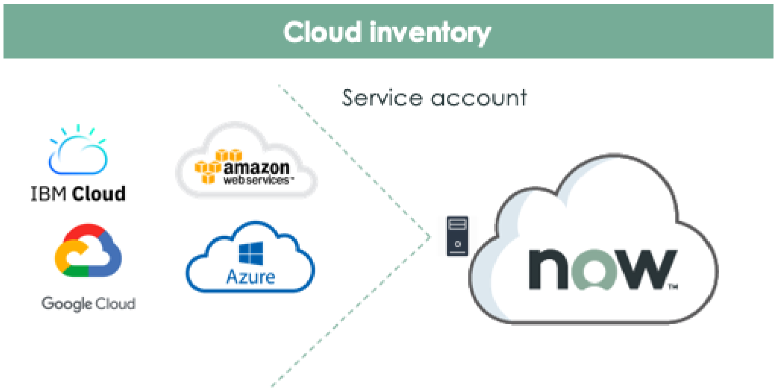

Visibility to IaaS/PaaS Services
Visibility to Tags

Service account-based Discovery leverages the REST APIs from major cloud providers to collect metadata on IaaS services, such as compute, network, and load balancer services. PaaS and FaaS services are also discovered with API-based Discovery.

NOTE: AWS and Azure cloud discovery use CAPI for IAAS discovery,[2] whereas IBM Cloud and Google Cloud Discovery use a pattern framework.

[1] https://www.wsj.com/articles/capital-one-breach-casts-shadow-over-cloud-security-11564516541

[2] https://docs.servicenow.com/bundle/newyork-it-operations-management/page/product/cloud-management-v2/concept/cloud-management-api.html

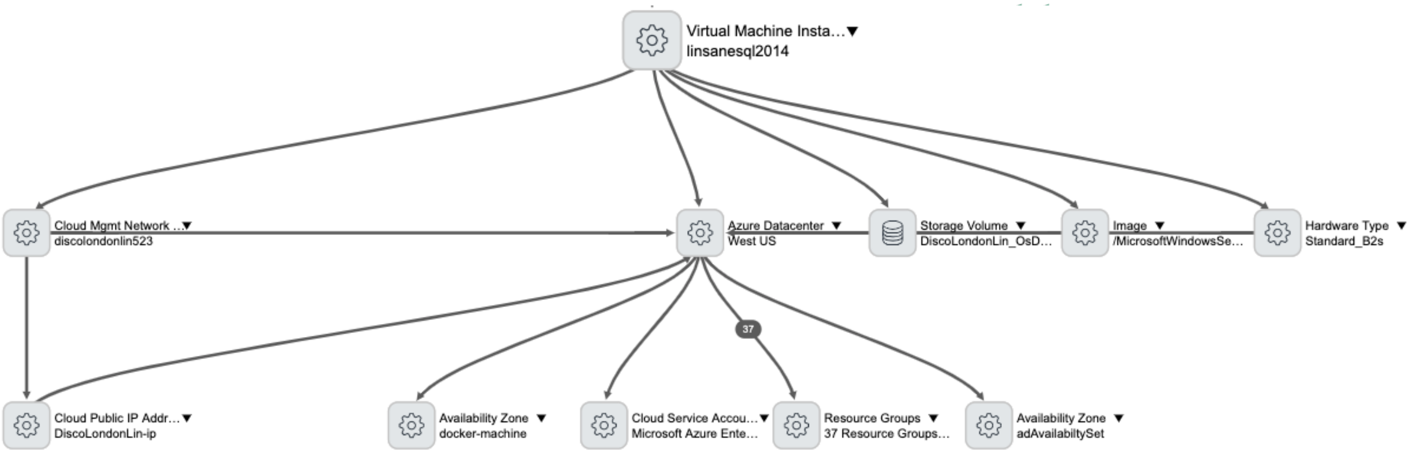

Metadata Discovery via API populates the Virtual Machine Instance CMDB table with basic attributes. For AWS Cloud, discovery framework support AWS organizations, where customer can add a master account and discovery can related member accounts which is part of the master account. For London, Madrid ServiceNow family release, discovery framework supports out-of-box IAM Role "OrganizationAccountAccessRole" [KB0725049]. With New York release, customers have option to configure custom roles in ServiceNow. For more details, check out the docs.servicenow.com. [Link]

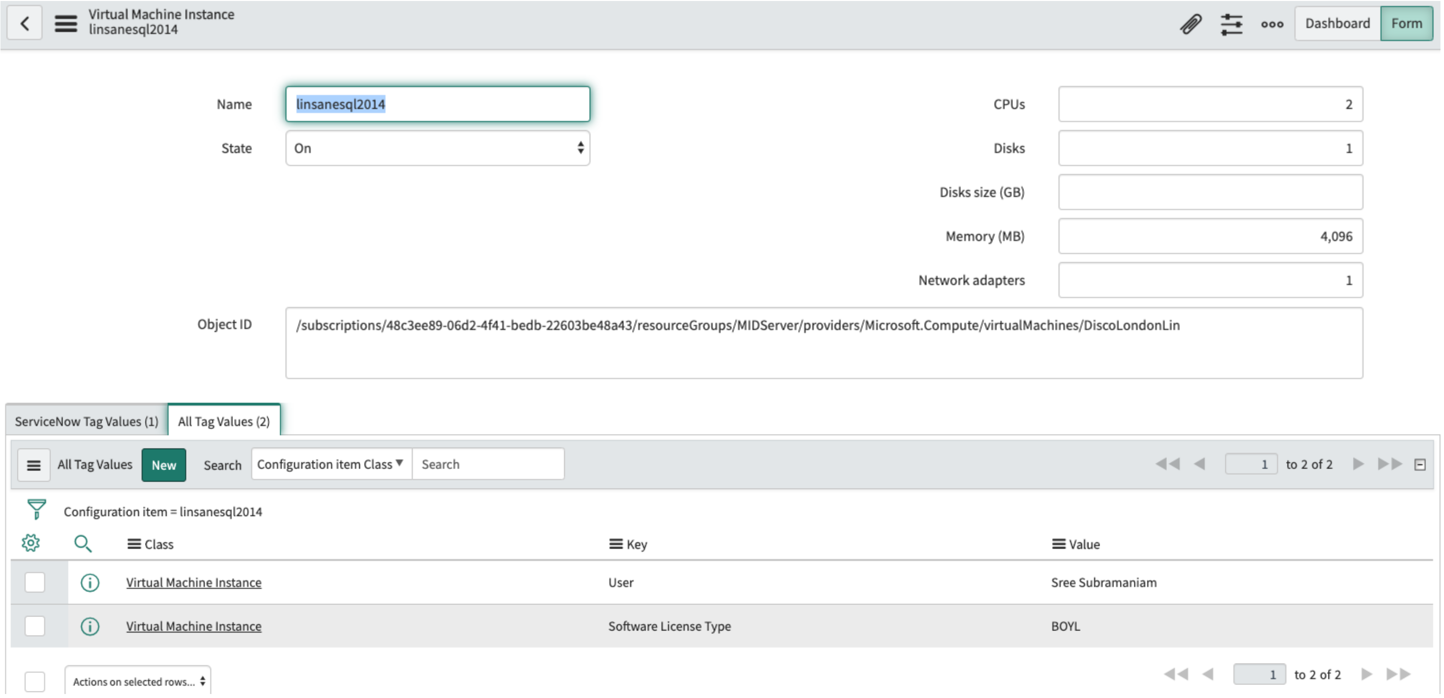

Virtual Machine Instance
linsanesql2014

Name	linsanesql2014	CPUs	2
State	On	Disks	1
		Disks size (GB)	
		Memory (MB)	4,096
		Network adapters	1
Object ID	/subscriptions/48c3ee89-06d2-4f41-bedb-22603be48a43/resourceGroups/MIDServer/providers/Microsoft.Compute/virtualMachines/DiscoLondonLin		

ServiceNow Tag Values (1) | All Tag Values (2)

All Tag Values | New | Search | Configuration item Class ▼ | Search | 1 to 2 of 2

Configuration item = linsanesql2014

Class	Key	Value
Virtual Machine Instance	User	Sree Subramaniam
Virtual Machine Instance	Software License Type	BOYL

Actions on selected rows... | 1 to 2 of 2

With the New York release, customers can use the Discovery Manager wizard to set up Cloud Discovery. This wizard enhances the Cloud Discovery experience and enables Discovery administrators to achieve their goals quickly. See Using the Discovery Manager for details.

Visibility into TAGS

Tags are labels in the form of key-value pairs that may be attached to cloud resources, such as instances, storage volumes, and databases. Tags provide additional information and context about a specific resource.

Service account-based Discovery collects tags associated with IaaS/PaaS components in CMDB key value pair tables. Real-time visibility into tag data in the CMDB helps with workflow automation and data reporting requirements use cases.

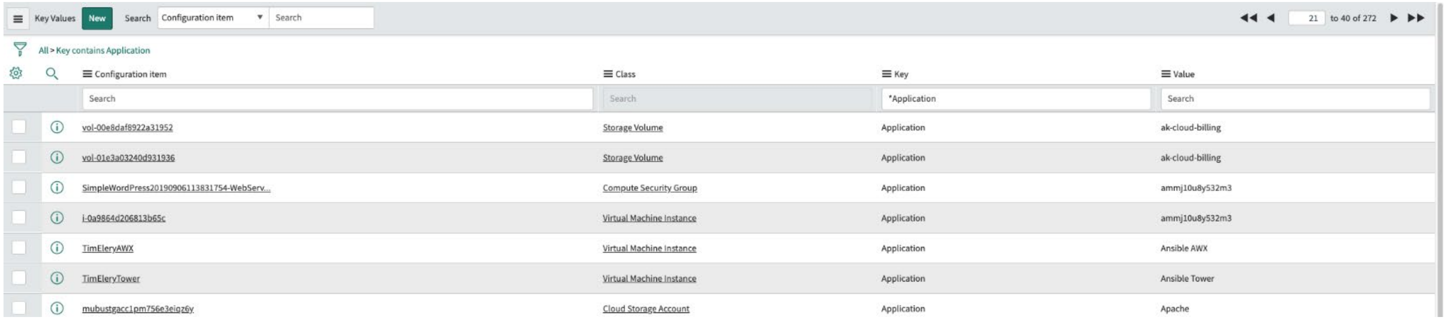

All > Key contains Application

Configuration item	Class	Key	Value
Search	Search	*Application	Search
vol-00e8daf8922a31952	Storage Volume	Application	ak-cloud-billing
vol-01e3a03240d931936	Storage Volume	Application	ak-cloud-billing
SimpleWordPress20190906113831754-WebServ...	Compute Security Group	Application	ammj10u8y532m3
i-0a9864d206813b65c	Virtual Machine Instance	Application	ammj10u8y532m3
TimEleryAWX	Virtual Machine Instance	Application	Ansible AWX
TimEleryTower	Virtual Machine Instance	Application	Ansible Tower
mubustgacc1pm756e3eiqz6y	Cloud Storage Account	Application	Apache

Deep-dive Discovery of cloud VMs

Metadata Discovery often helps customers with basic data visibility on VMs. However, the data is not usable for outcomes like Software Asset Management or AIOps or SecOps use cases. The compute layer should be interrogated by IP-based discovery schedules to collect enriched data sets, such as Installed software data, process information, TCP/IP connections, enterprise applications deployed on cloud VMs, and so on.

IP-based discovery to enrich meta
data from service account discovery.

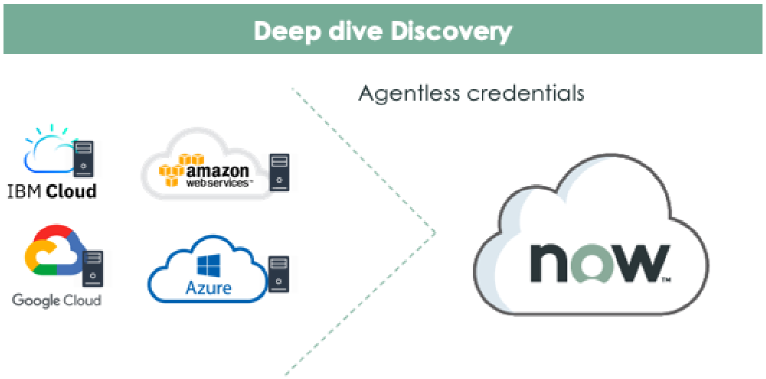

Deploy MIDs in Transit VPC/Management
VPCs and get end-to-end visibility to
VMs running on public cloud

For deep-dive discoveries, customers should deploy MID Servers in a Virtual Private Cloud (VPC), and then activate IP-based discovery jobs to access the compute layer using agentless credentials.

Consider the example below, in which a customer has deployed an Oracle Database on an EC2 compute. Metadata Discovery provides only the image of the virtual machine instance relationships. For example, deep level interrogations with EC2 compute is required to collect data, such as the Oracle Database name, version, edition, oracle options, and oracle catalog data required for Software Asset Management, AIOps, and SecOps use cases.

> For deep-dive discoveries, customers should deploy MID Servers in a Virtual Private Cloud (VPC), and then activate IP-based discovery jobs to access the compute layer using agentless credentials.

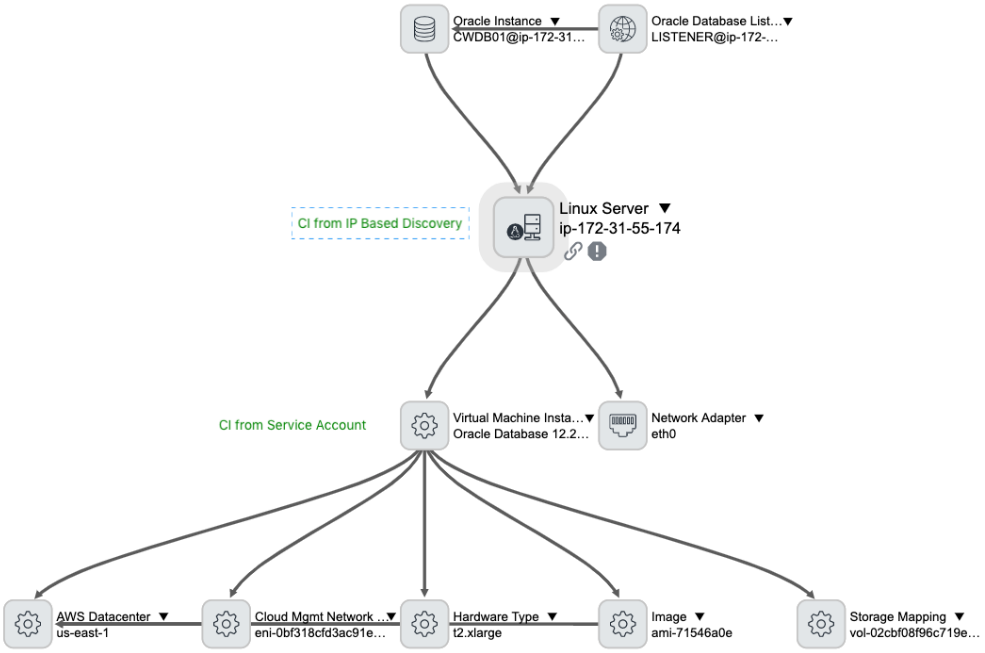

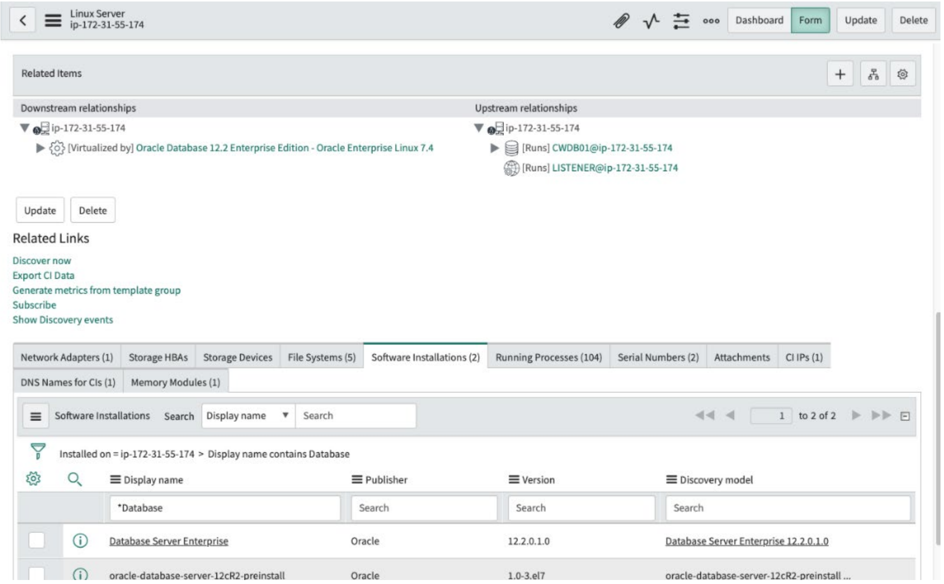

Related Items

Downstream relationships	Upstream relationships
ip-172-31-55-174	ip-172-31-55-174
[Virtualized by] Oracle Database 12.2 Enterprise Edition - Oracle Enterprise Linux 7.4	[Runs] CWDB01@ip-172-31-55-174
	[Runs] LISTENER@ip-172-31-55-174

Related Links

Discover now
Export CI Data
Generate metrics from template group
Subscribe
Show Discovery events

Installed on = ip-172-31-55-174 > Display name contains Database

Display name	Publisher	Version	Discovery model
*Database	Search	Search	Search
Database Server Enterprise	Oracle	12.2.0.1.0	Database Server Enterprise 12.2.0.1.0
oracle-database-server-12cR2-preinstall	Oracle	1.0-3.el7	oracle-database-server-12cR2-preinstall ...

Customers with multiple VPC segments can deploy the MID servers in a management or transit VPC. For example, this best practice recommendation from AWS can be adopted for MID Server deployments. [Reference[3]]

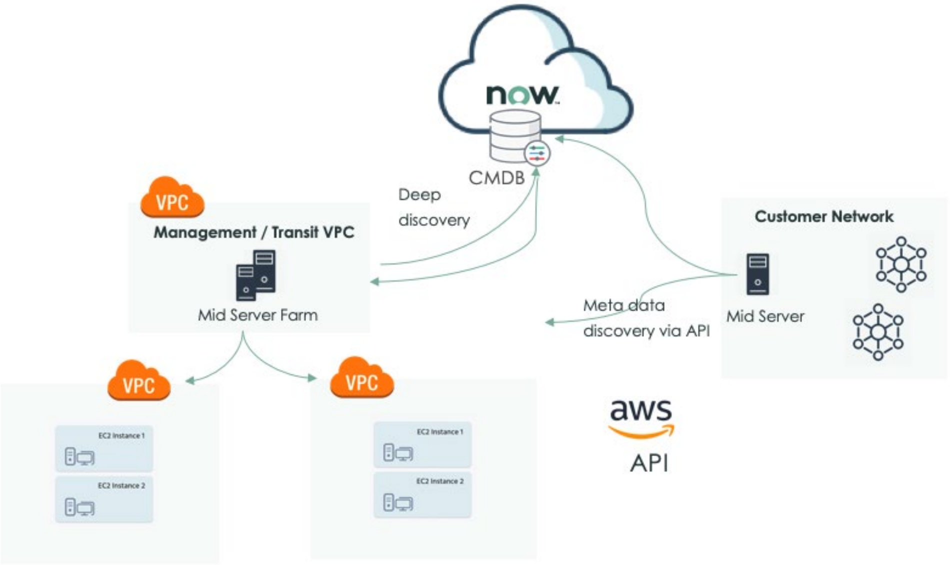

Event-driven Discovery

The majority of cloud workloads are ephemeral in nature. It is imperative to have the CMDB reflect the near real-time nature of cloud workloads to effectively automate workflows. ServiceNow Discovery supports event-driven discovery. Customers can push cloud events to ServiceNow, which triggers Discovery on demand on the target service to identify changes and update the CMDB in near real time.

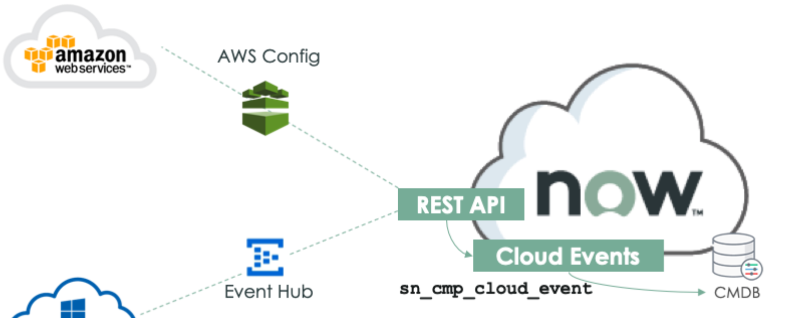

Get near real-time updates to CMDB

See this documentation for details:

- Configure the Amazon AWS Config service to auto-update the CMDB
- Configure the Azure Alert service to auto-update the CMDB

> The majority of cloud workloads are ephemeral in nature. It is imperative to have the CMDB reflect the near real-time nature of cloud workloads to effectively automate workflows. ServiceNow Discovery supports event-driven discovery.

[3] https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/transit-vpc.html

Discovery serverless workloads

Serverless computing is a cloud computing model which aims to abstract server management and low-level infrastructure decisions away from developers.

AWS was first to enter the serverless market in 2014, and its Lambda platform continues to be synonymous with the concept of serverless computing. It wasn't until 2016 that Google introduced Cloud Functions. Shortly after that, Microsoft released Azure Functions. It's been over 3 years since functions were introduced, and now each of the major cloud providers offer compelling services. However, serverless computing is still a buzzword, and the ecosystem is fairly immature.

The ServiceNow Discovery framework fully supports the discovery of serverless functions like Lambda and Azure. Solutions also include Service Mapping capabilities for cloud-native serverless components. Discovery patterns can detect functions and populate the **cmdb_ci_cloud_function** table.

> “
>
> The ServiceNow Discovery framework fully supports the discovery of serverless functions like Lambda and Azure.

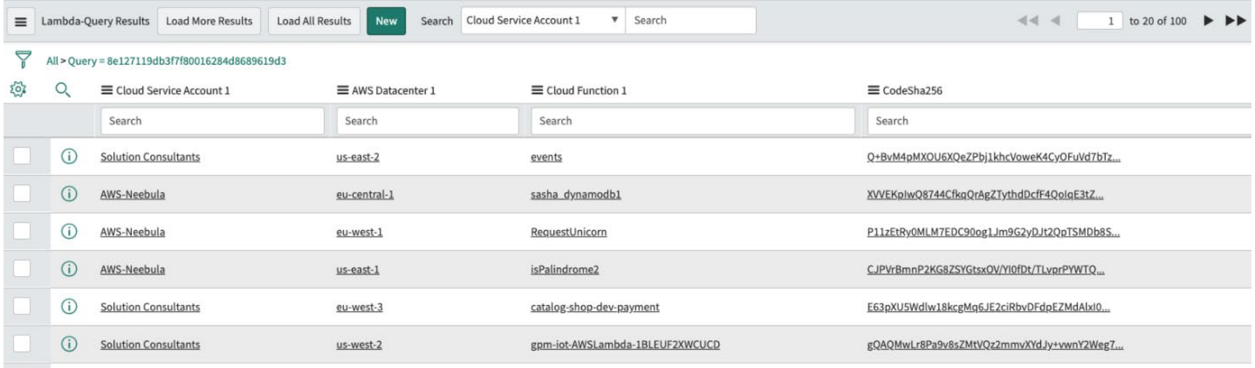

All > Query = 8e127119db3f7f80016284d8689619d3

Cloud Service Account 1	AWS Datacenter 1	Cloud Function 1	CodeSha256
Solution Consultants	us-east-2	events	Q+BvM4pMXOU6XQeZPbj1khcVoweK4CyOFuVd7bTz...
AWS-Neebula	eu-central-1	sasha_dynamodb1	XVVEKplwQ8744CfkqQrAgZTythdDcfF4QolqE3tZ...
AWS-Neebula	eu-west-1	RequestUnicorn	P11zEtRy0MLM7EDC90og1Jm9G2yDJt2QpTSMDb8S...
AWS-Neebula	us-east-1	isPalindrome2	CJPVrBmnP2KG8ZSYGtsxOV/YI0fDt/TLvprPYWTQ...
Solution Consultants	eu-west-3	catalog-shop-dev-payment	E63pXU5Wdlw18kcgMq6JE2ciRbvDFdpEZMdAlxI0...
Solution Consultants	us-west-2	gpm-iot-AWSLambda-1BLEUF2XWCUCD	gQAQMwLr8Pa9v8sZMtVQz2mmvXYdJy+vwnY2Weg7...

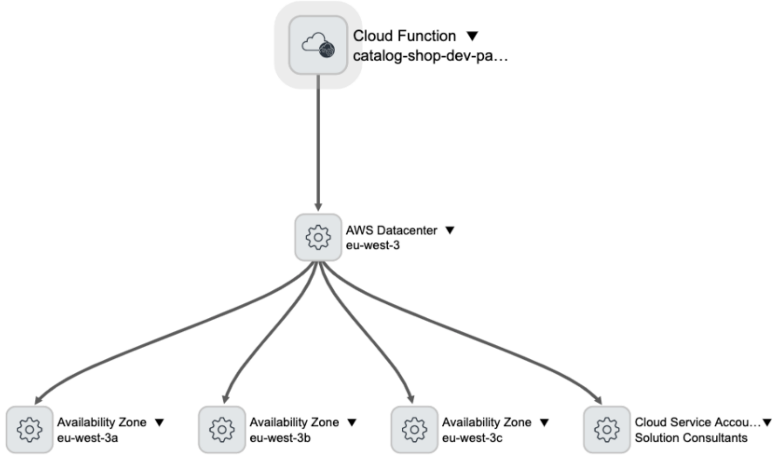

Lambda and Azure functions can be added as entry points for Service Mapping which allow Discovery to detect function-to-function calls. In the example on the next page, the service leverages lambda functions, an S3 bucket, and an API gateway as the technology stack.

An HTTPS entry point for top-down mapping can detect function-to-function calls and can provide a holistic view of the application dependency view on the cloud-native application. Maps created by Service Mapping play a vital role in change impact analysis. These maps provide impact visualization and enable topology-based event correlation for AIOps.

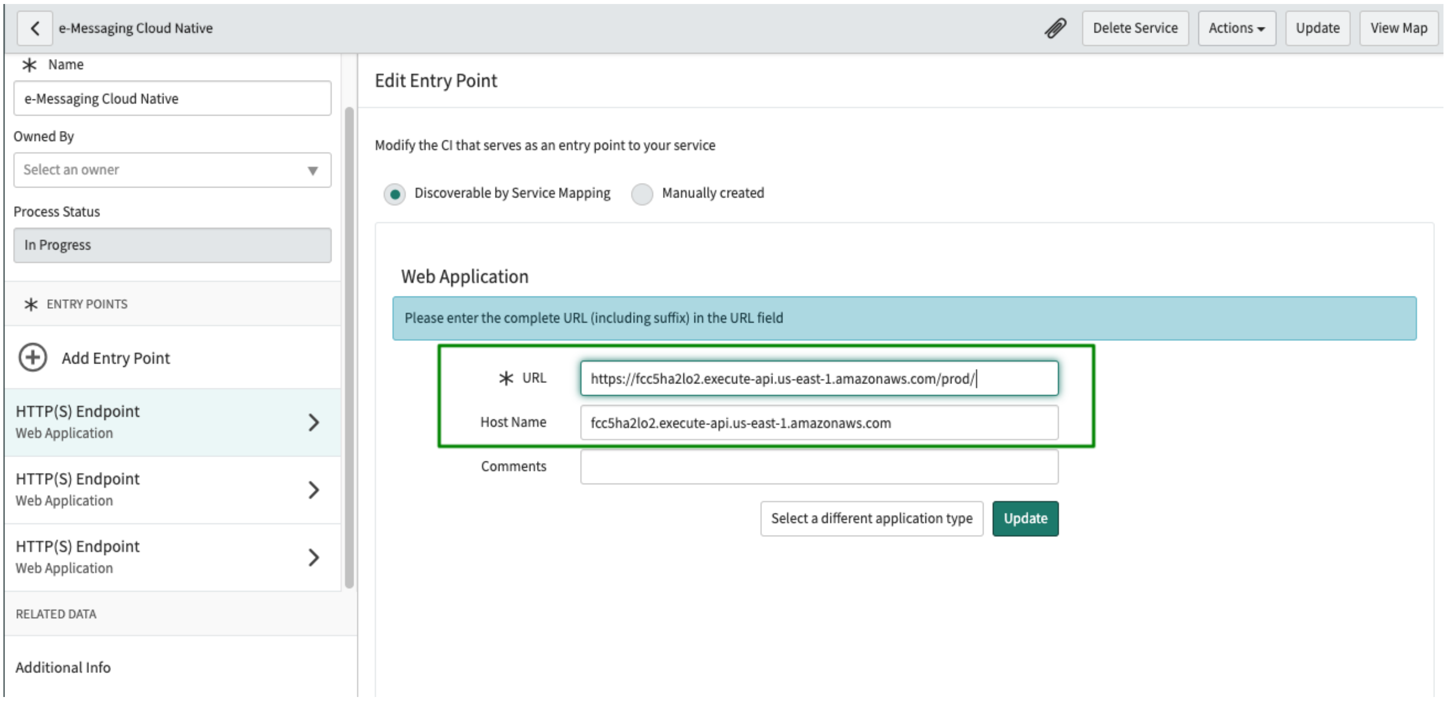

e-Messaging Cloud Native

✱ Name: e-Messaging Cloud Native

Owned By: Select an owner

Process Status: In Progress

HTTP(S) Endpoint
Web Application

HTTP(S) Endpoint
Web Application

HTTP(S) Endpoint
Web Application

Edit Entry Point

Modify the CI that serves as an entry point to your service

● Discoverable by Service Mapping ○ Manually created

Web Application

Please enter the complete URL (including suffix) in the URL field

✱ URL: https://fcc5ha2lo2.execute-api.us-east-1.amazonaws.com/prod/

Host Name: fcc5ha2lo2.execute-api.us-east-1.amazonaws.com

Comments:

This is an AIOps view of a cloud-native app with Lambda function.

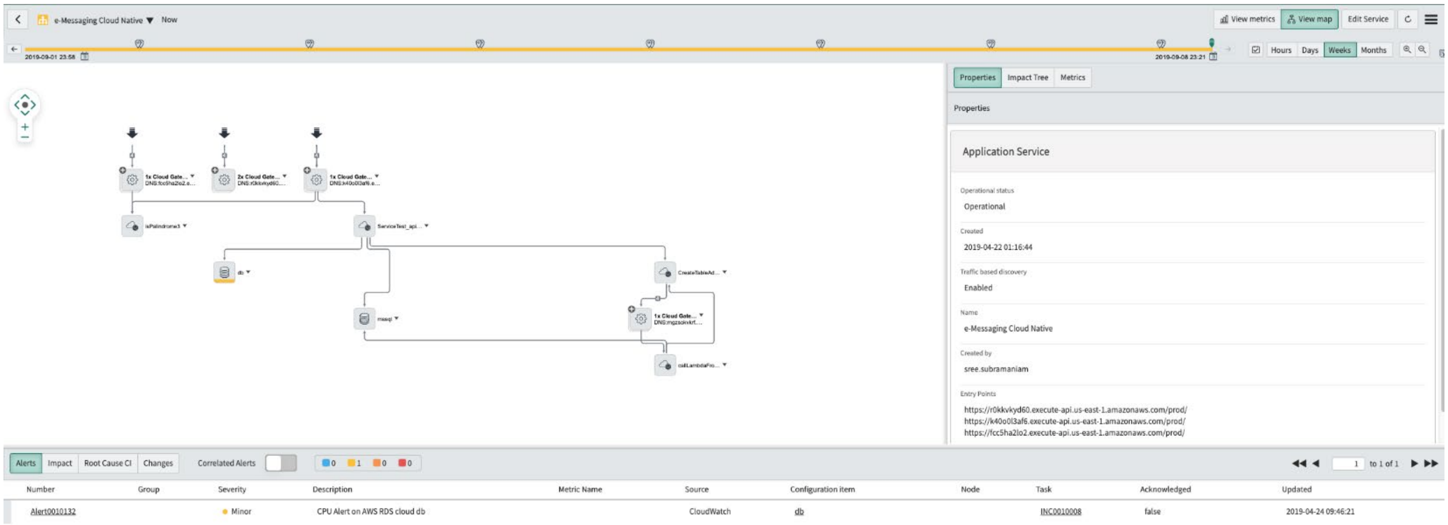

Properties

Application Service

Operational status: Operational

Created: 2019-04-22 01:16:44

Traffic based discovery: Enabled

Name: e-Messaging Cloud Native

Created by: sree.subramaniam

Entry Points:
https://r0kkvkyd60.execute-api.us-east-1.amazonaws.com/prod/
https://k40o0l3af6.execute-api.us-east-1.amazonaws.com/prod/
https://fcc5ha2lo2.execute-api.us-east-1.amazonaws.com/prod/

Number	Group	Severity	Description	Metric Name	Source	Configuration item	Node	Task	Acknowledged	Updated
Alert0010132		● Minor	CPU Alert on AWS RDS cloud db		CloudWatch	db		INC0010008	false	2019-04-24 09:46:21

New York Release enhancements

- **AWS—Identity and Access Management (IAM) roles**
  Customers can configure an IAM role to provide temporary security credentials that a MID Server can use to discover cloud resources. For details, see Configure the MID Server for AWS IAM roles.
- **AssumeRole enhancements for AWS organizations**
  Support for AWS Organizations, previously introduced in the London release, now leverages fully configurable AssumeRole request parameters as dictated by the AWS Security Token Service AssumeRole API Action.

For details on the specific parameters, see these documents:

- AWS Documentation for the AssumeRole API Action
- Assuming member roles with an AWS API
- AWS Organizations discovery is not finding cloud resources

> “
> Setting up Cloud Discovery requires only a few clicks in the New York release.

Cloud Discovery UI improvements

Setting up Cloud Discovery requires only a few clicks in the New York release. Use these steps to set up Cloud Discovery from the Discovery Manager wizard. See Using the Discovery Manager for details.

Step 1: Navigate to **Discovery > Home** and add a cloud schedule.

Step 2: Choose the cloud provider, create a unique schedule name, add or choose service account credentials, and then test access to the account.

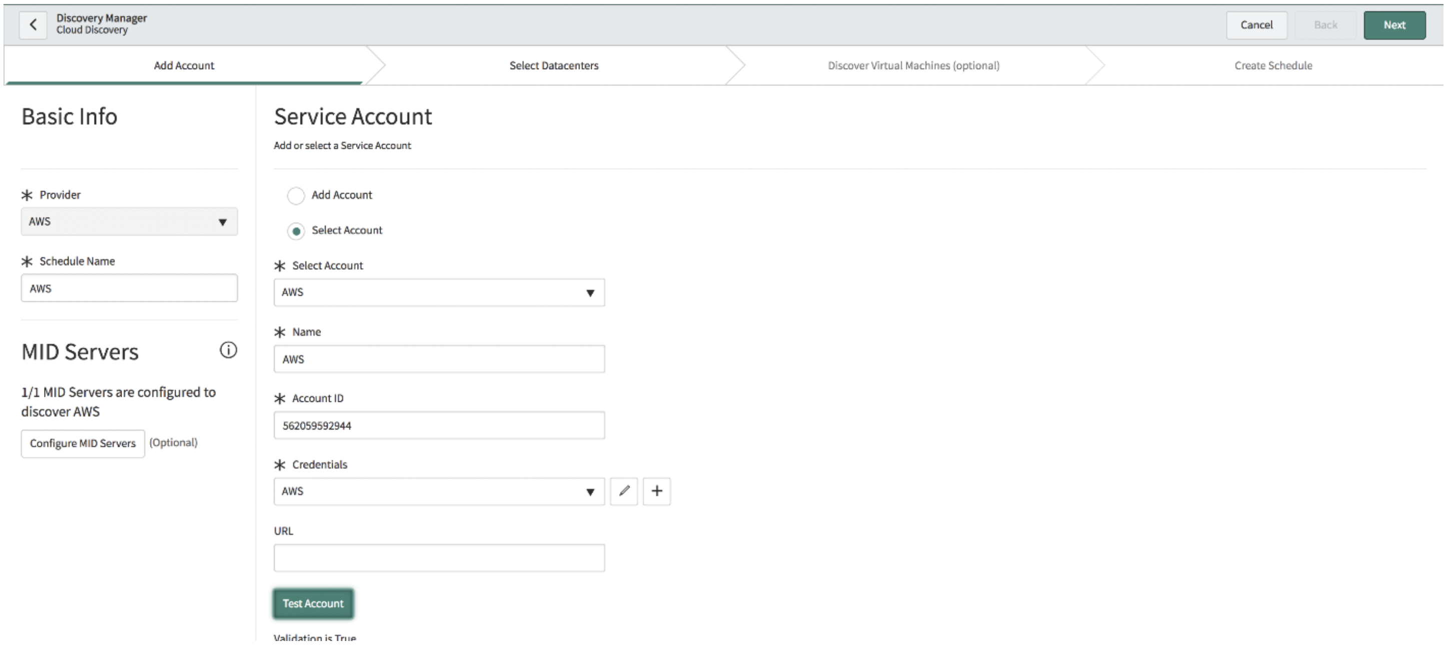

Discovery Manager
Cloud Discovery

Cancel | Back | Next

Add Account | Select Datacenters | Discover Virtual Machines (optional) | Create Schedule

Basic Info

* Provider: AWS

* Schedule Name: AWS

MID Servers

1/1 MID Servers are configured to discover AWS

Configure MID Servers (Optional)

Service Account

Add or select a Service Account

○ Add Account

● Select Account

* Select Account: AWS

* Name: AWS

* Account ID: 562059592944

* Credentials: AWS

URL:

Test Account

Validation is True

Step 3: Choose one or more cloud provider datacenters (regions) to discover.

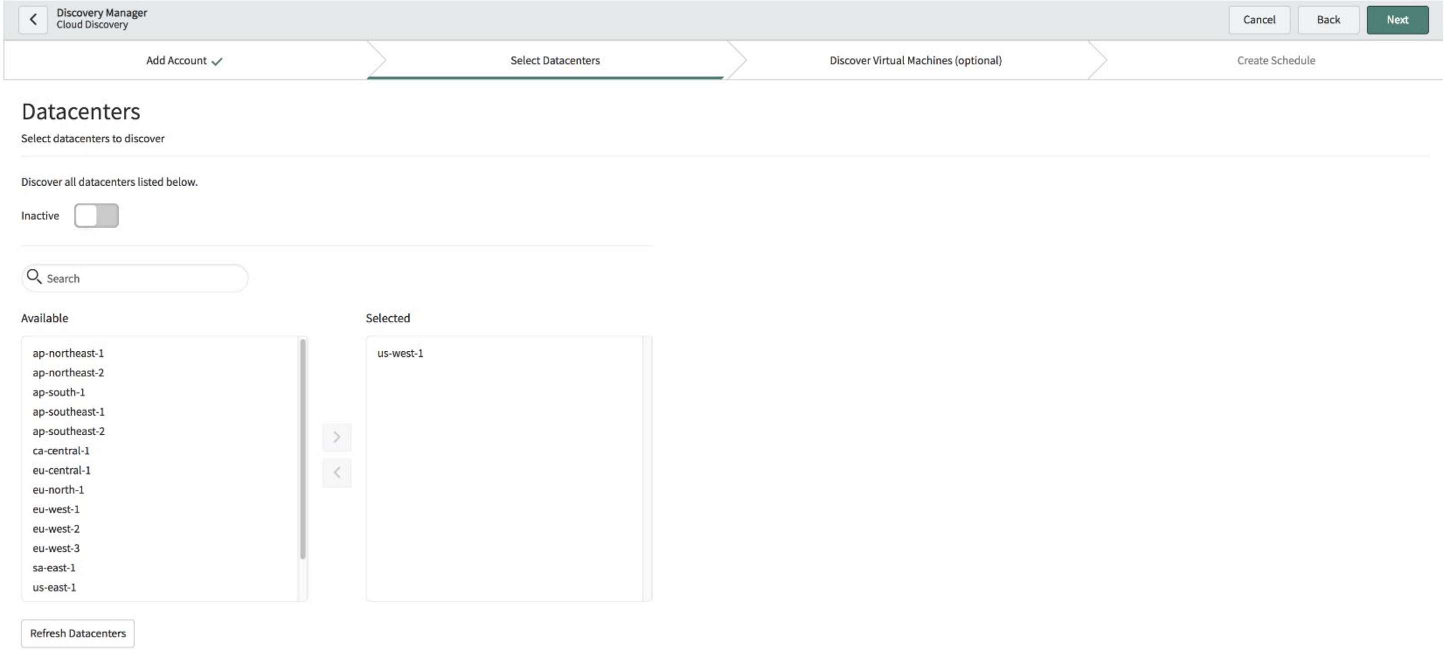

Datacenters

Select datacenters to discover

Discover all datacenters listed below.

Inactive

Available

ap-northeast-1
ap-northeast-2
ap-south-1
ap-southeast-1
ap-southeast-2
ca-central-1
eu-central-1
eu-north-1
eu-west-1
eu-west-2
eu-west-3
sa-east-1
us-east-1

Selected

us-west-1

Step 4: Create a Discovery schedule and choose the scanning frequency.

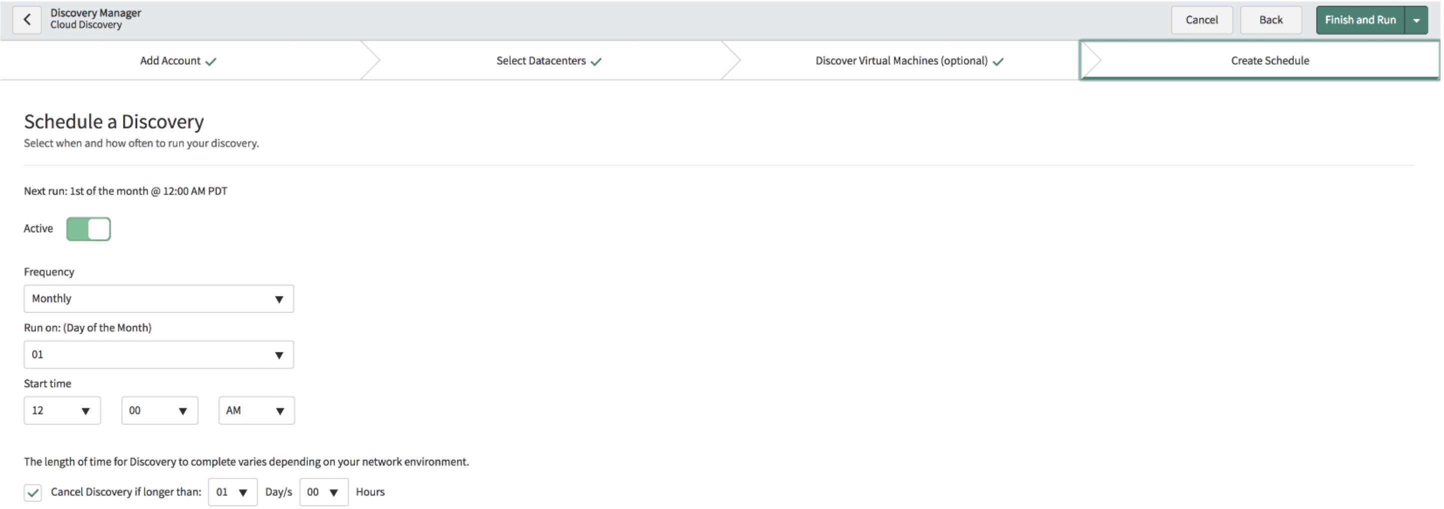

Schedule a Discovery

Select when and how often to run your discovery.

Next run: 1st of the month @ 12:00 AM PDT

Active

Frequency: Monthly

Run on: (Day of the Month): 01

Start time: 12 : 00 : AM

The length of time for Discovery to complete varies depending on your network environment.

☑ Cancel Discovery if longer than: 01 Day/s 00 Hours

Step 5: Visualize the results.

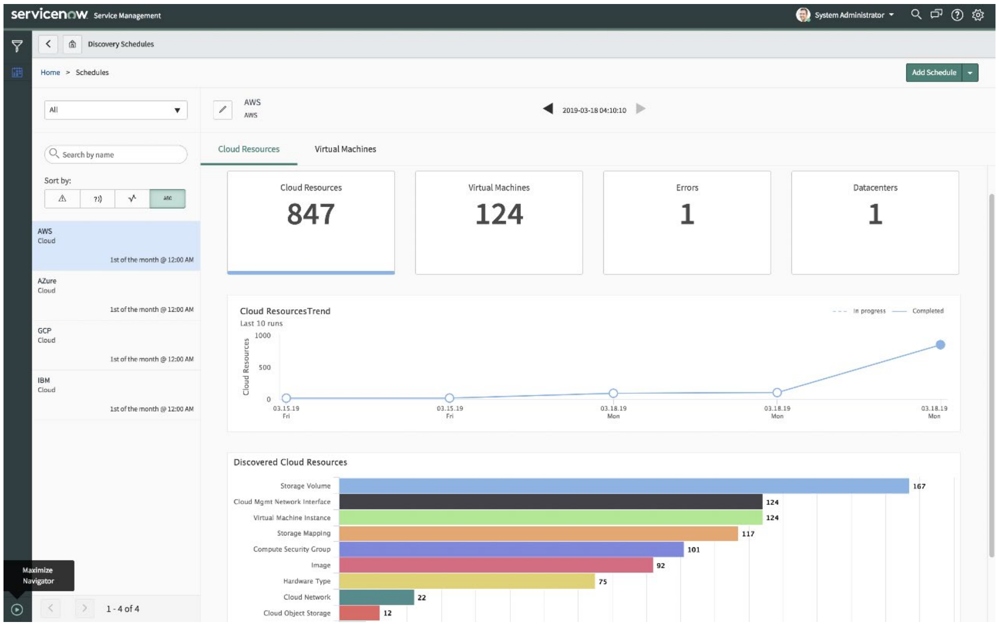

AWS
Cloud
1st of the month @ 12:00 AM

AZure
Cloud
1st of the month @ 12:00 AM

GCP
Cloud
1st of the month @ 12:00 AM

IBM
Cloud
1st of the month @ 12:00 AM

1 - 4 of 4

AWS
AWS

2019-03-18 04:10:10

Cloud Resources | Virtual Machines

Cloud Resources	Virtual Machines	Errors	Datacenters
847	124	1	1

ITOM Visibility low-code framework—pattern engine and out-of-band patterns

The ITOM Visibility pattern framework provides a low-code SDK that the ServiceNow R&D team uses to build and ship Discovery patterns. This framework also allows customers to personalize and extend out-of-box patterns.

The ServiceNow product team decoupled the pattern releases from the semi-annual platform releases and now offers patterns on store.servicenow.com. Customers can Opt-In for the Store application to get monthly updates on new pattern content.

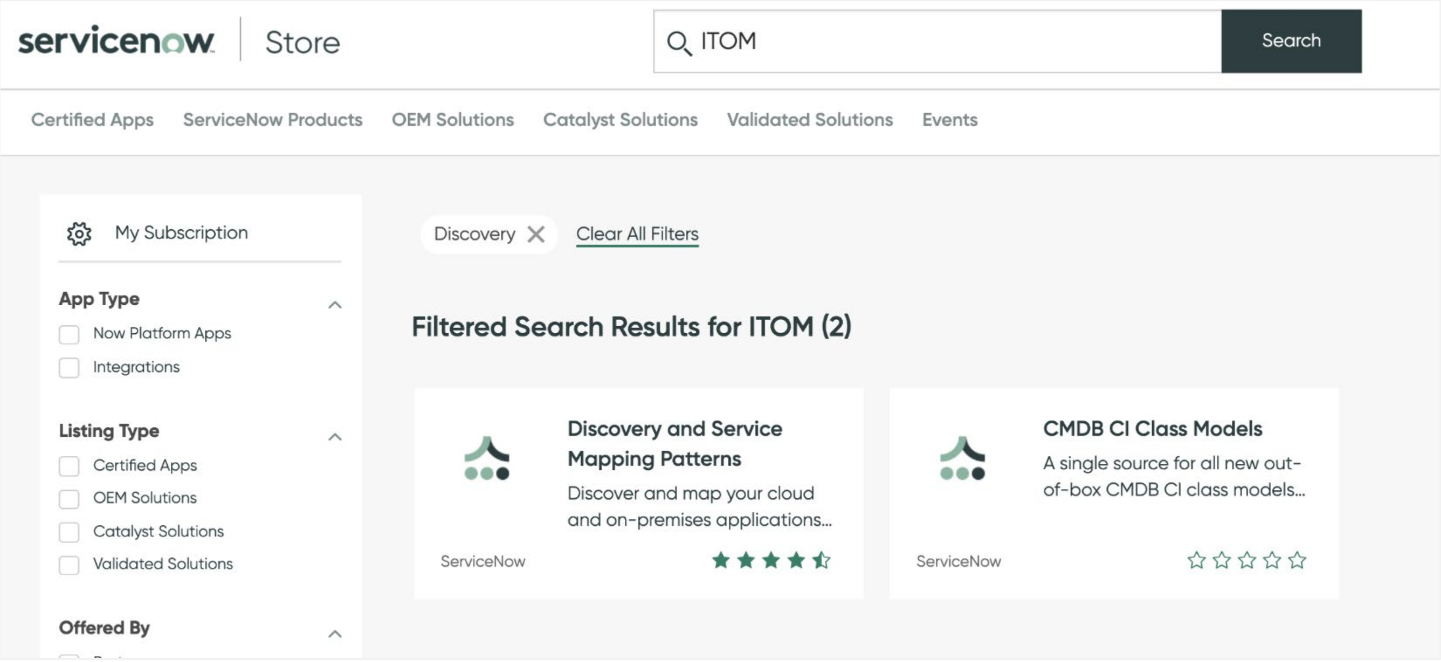

ITOM

□ Now Platform Apps
□ Integrations

□ Certified Apps
□ OEM Solutions
□ Catalyst Solutions
□ Validated Solutions

Discovery ✕

Filtered Search Results for ITOM (2)

Discovery and Service Mapping Patterns
Discover and map your cloud and on-premises applications...
ServiceNow ★★★★⯪

CMDB CI Class Models
A single source for all new out-of-box CMDB CI class models...
ServiceNow ☆☆☆☆☆

Some of the top Cloud Discovery patterns available from Store:

- Amazon DynamoDB discovery
- AWS API gateway discovery
- AWS Cognito discovery
- AWS Lambda discovery
- AWS S3 discovery
- AWS tag discovery
- Google Cloud Platform discovery
- IBM Cloud Platform discovery
- Microsoft Azure Application Gateway discovery
- Microsoft Azure resource discovery
- Microsoft Azure Functions discovery

For more information:

- https://www.servicenow.com/products/discovery.html
- https://www.servicenow.com/products/service-mapping.html

servicenow

SN-Whitepaper-VisibilitytoMulti-Cloud-102119