ivanti

Neurons for Service Mapping-Discovery Requirements

Ver. 3.3.0

Technical Document

Table of Contents

Introduction

About This Document

This document provides much of the information needed to properly prepare the discovery portion of Ivanti Neurons for Service Mapping. It also provides a check list of the system requirements to run the Discovery Application ("Discovery App") on customer-hosted servers, the firewall rules and ports to allow communication between the Discovery App and cloud instance, and the credentials needed to accurately identify infrastructure assets, gather configurations, and detect application dependencies and inter-relationships. It is not a how-to-guide for implementing, configuring, or troubleshooting Ivanti Neurons for Service Mapping.

About Neurons for Service Mapping

Ivanti Neurons for Service Mapping provides Ivanti Neurons for ITSM customers with automatic inventory and configuration cataloging of on-premise and cloud compute, network and storage assets. The discovery process also detects physical and logical dependencies between applications and their host systems. From there, enterprise services are easily mapped across all supporting infrastructure components and viewable from within the Ivanti CMDB. Neurons for Service Mapping offers the ease of SaaS deployment and accessibility combined with the security and performance of on-premise discovery. Leveraging over 140 (and growing) agentless and extendable probes.

Discovery Architecture Overview

Figure 1 provides a visual representation of the methods Neurons for Service Mapping employs to perform agentless asset, configuration, software and relationship/dependency discovery.

Not depicted in Figure 1 is the option of deploying agents to Windows and Linux servers in lieu of agentless scanning methods. These are the options for agent to Windows and Linux servers:

- Windows EXE and Linux TAR files can be downloaded via the Service Mapping UI to be manually installed on target host systems or distributed via customer's software distribution tools
- Windows Agent can be pushed to target hosts via the Discovery Application leveraging the same Batch requirements for agentless scanning (see Figures 5 and 6)
- Ivanti can provide an MSI build of the agent (development fees apply)

Figure 1: Agentless Discovery Overview

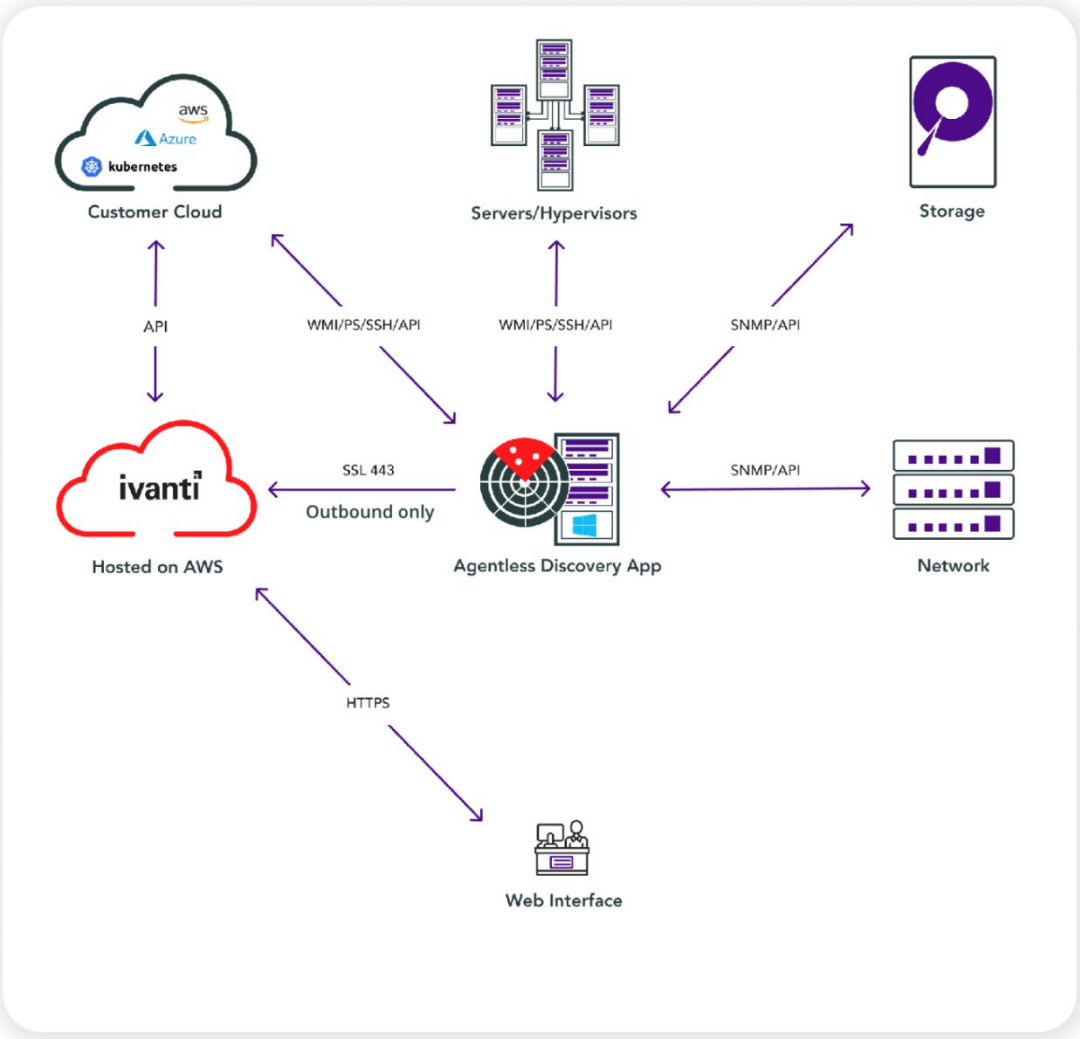

Discovery App Sizing Considerations

The number and locations of Discovery Apps depend on several factors. These include size of network to scan, number of discoverable assets, time required to complete all scans, and number of logically or physically separated networks. Here are some questions to consider:

- How many data centers or cloud environments to be scanned?
- What is the connectivity/bandwidth between each data center?
- Would any firewall rules/router settings prevent scanning between data centers? (see authentication requirements for necessary ports)
- How many IP subnets per DC/cloud environment?

- How many devices/virtual servers in total?
- Roughly how many devices/virtual servers per subnet?
- What is the desired scan frequency (i.e. daily, weekly, alternating)? Any exceptions or special considerations?

Discovery App Scan Capabilities

Each Discovery App is capable of simultaneously scanning subnets of four /24, two /23 or one /22. Scan times can vary based on the number and types of connected devices. Generally, a scan of a /24 subnet with the maximum number of 256 devices will take 20-40 minutes to complete. This means four /24 subnets can be simultaneously scanned in an average time of 30 minutes. Figure 2 shows the estimated scan times for various amounts of /24 subnets and Discovery Apps.

Figure 2: Estimated Scan Times

Number of /24 Subnets	Number of IPs	Number of Discovery Apps	Hours to complete scan (assumes 4 simultaneous subnets)
4	1,024	1	.5
25	6,400	1	3
100	25,600	1	13
200	52,200	2	13
400	102,400	4	13

As you can see, a single Discovery App can easily scan over 20,000 IPs in a single day. If daily scans for each IP are not required, subnet scans can be configured to run on alternating days so even more IP's can be scanned from a single Discovery App each week. If, however, the desired scan frequency is greater than once a week per IP, and the subnet count is higher than 100, more than one Discovery App may be required.

Discovery App Server Requirements

The Discovery App must be installed on a Windows server (see Figure 3 for recommended specs) regardless of the types of devices or operating systems to be discovered. The Windows server must be located within

the domain to be discovered or trust across domains must be enabled. The Discovery App Windows server must also have a persistent outbound connection to the URLs shown in Figure 4.

The Discovery App will be downloaded from a link inside the Virima UI. Please be sure to have a browser installed (Chrome recommend) and access the region specific URL listed in Figure 4.

Figure 3: Discovery App System Requirements

Operating System	Server Type	CPU	RAM	HD	Browser
Windows Server 2012 or newer	Dedicate Physical or Virtual	High performance (12-16 cores)	16 GB	50 GB	Chrome

Firewall Rules

Discovery is performed when the Discovery App launches selective probes to the target systems. These scans are configured via the cloud-based UI which in turn commands the Discovery App to launch the probes. The data collected by the probes is then returned to the Discovery App which is then securely transmitted to the Ivanti cloud for processing and staging for synch into the Ivanti CMDB. Figure 4 shows the necessary communication that must be allowed between the Discovery App(s) located in the customer's on-premise and/or cloud environments and the Ivanti cloud.

Figure 4: Firewall Rule Requirements

<table>
<tr><th colspan="2">Discovery App communication to INSM cloud requires 443 to the following region-specific URLs</th></tr>
<tr><td>U.S. Hosting
https://discoveryselector-usw2.ivanticloud.com
https://discoveryserver1-usw2.ivanticloud.com
https://servicemapping-usw2.ivanticloud.com</td><td>EU (Frankfurt) Hosting
https://discoveryserver1-euc1.ivanticloud.com
https://discoveryselector-euc1.ivanticloud.com</td></tr>
<tr><td>Australia Hosting
https://discoveryserver1-apse2.ivanticloud.com/
https://discoveryselector-apse2.ivanticloud.com/</td><td></td></tr>
<tr><th colspan="2">Proxy Server</th></tr>
<tr><td colspan="2">Proxy server is supported* for outbound communication from Discovery App to INSM cloud with the following properties in Discovery Application common.properties file:
#ProxyHost= #ProxyPort= #ProxyUserName= #ProxyPassWord=

*Currently a proxy server cannot be used in conjunction with CyberArk PAM integration</td></tr>
<tr><th colspan="2">Allow Persistent Connection</th></tr>
</table>

Firewall rules must allow for a persistent connection between the Discovery Application and Service Mapping cloud. Same is true for connections between Discovery Agents and the Discovery Application.
ICMP Ping or NMAP
ICMP Ping is the default mechanism to perform host up/down check of target devices at the start of every scan. If Ping is not allowed, the NMAP option needs to be enabled in Discovery Admin settings-> Discovery Scan Configuration-> "Enable port check using NMAP"

Port Requirements

Figure 5 details the ports required for the Discovery App to conduct agentless scanning, receive data from the Windows agents, and transfer the discovery data to the Neurons for Service Mapping cloud (depicted as "Discovery Server").

Figure 5: Port Requirements

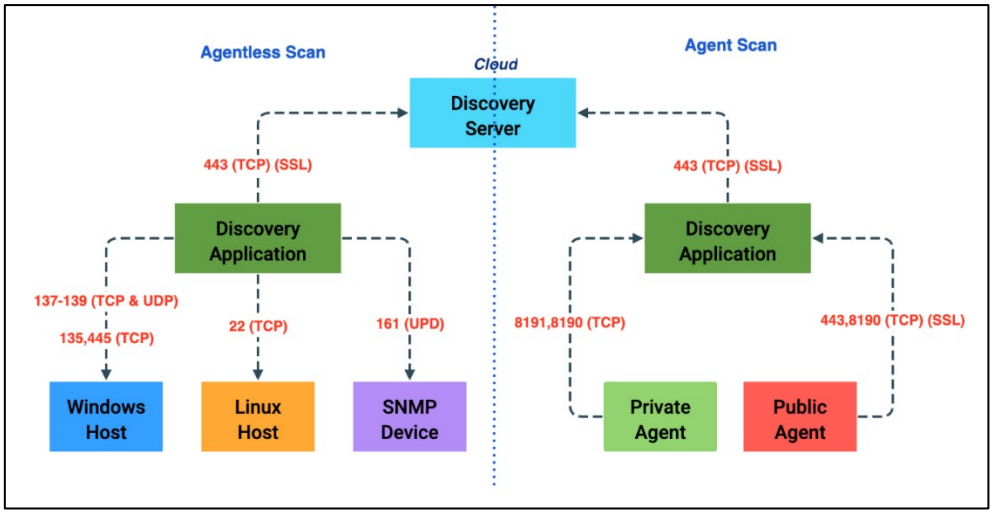

Authentication Requirements

As part of the Discovery App installation, customers enter the appropriate credentials for the types of systems to be scanned. These credentials are stored and encrypted locally within the DA Windows server and never go to the cloud. Figure 6 lists the most common types of

credentials and communication methods required for the probes to successfully scan each asset type. Note the list may not be inclusive of requirements and is subject to change.

Notes on scanning Windows OS via Batch (WMI and PAExec): The Discovery App will copy a few files to the Windows admin$ share and then launch a local service with the credentials provided. It will then scan the system with all the probe details. Once the scan is complete it will then stop and remove the service and files from the admin$ share folder. In order for this to work, the account being used must have remote file access to the admin$ share, must have admin rights on the box being scanned, and must have logon as a service rights on the local Windows box.

Figure 6: Discovery Credential and Port Requirements

Asset Type	Method	Port/Protocol	Authentication Method	Notes
Windows (Batch method)	Batch scripts with WMI commands and PAExec from Disc. App.	TCP: 135, 139 (WMI) and 445 (File and Print sharing) UDP: 137,138	Windows admin credentials with elevated privileges	Elevated privileges required for relationship discovery WMI and PAExec is required to push Windows Discovery Agents via the Discovery App.
Windows (PS remoting method)	PowerShell from Disc. App.	Port 5985	Windows username/password	PS-Remoting must be enabled on target hosts and Discovery App
Unix, Linux, Mac, Solaris	SSH from Disc. App.	22 (SSH)	SSH username/ password or private keys	Requires SUDO on netstat or SS command to fetch IP connections. SUDO dmidecode required to fetch serial number. SNMP will not discover relationships.
Network Infrastructure (via SSH)	SSH (Cisco CDP) from Disc. App.	22 (SSH)	SSH User ID/password	SSH required to discover relationships.
SNMP (Network Infrastructure, Windows, Linux)	SNMP from Disc. App.	161 (SNMP)	Community String (v1, v2) or User ID/password (v3)	SNMP does not provide sufficient capability to discover relationships for any device type so other methods are preferred.
AWS resources	API from cloud UI	443/80	Account ID, AWS access key, secret key, account ID	See *AWS IAM Policy* below, EC2 OS details require server discovery scan
Azure resources	API from cloud UI	443/80	Account ID, client ID, tenant ID, secret key	See *Azure Role Setup* below, Virtual machine details require server discovery scan
VMware vCenter	API from Disc. App.	443/80	VCenter SSO User ID and password	This will discover ESXi hosts managed by vCenter

ESXi Hosts	SSH from Disc. App.	22 (SSH)	SSH Host Credentials	For standalone ESXi host discovery
Active Directory	WMI/API	135, 88	AD host, Domain, Base DN, Bind DN, Password	Domain admin required
MS SQL Database	WMI from Disc. App.	1433, 1434	Windows or SQL Server account	Requires SQL server credential profile with user read rights on the sys.databases table
MMC Certificates	WMI and PowerShell from Disc. App.	135, 137, 138, 139 (WMI) and 445 (file/print sharing)	Windows admin credentials with elevated privileges	
TLS/SSL Certificates	WMI from Disc. App.	443	Windows admin credentials with elevated privileges	Certs that reside on a server can be discovered. E.g. if port 443 is open on a server, it is assumed it contains a SSL cert which Virima tries to retrieve from the port query
IIS Websites	WMI from Disc. App.	135, 137, 138, 139 (WMI) and 445 (file/print sharing)	Webserver credentials	Requires IIS Management Tools installed with IIS Management Scripts and Tools selected
Cisco Meraki	API from cloud UI	443/80	API key	Requires Meraki Cloud Dashboard API access

AWS IAM Policy

To allow the importing of AWS objects into the CMDB the following must be done within the AWS account. Note: EC2 operating system discovery is performed via the Discovery App per the Windows/Linux/Unix methods listed in the Credential Requirements Table above.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "ecr:Describe*",
        "ec2:Describe*",
        "rds:Describe*",
        "elasticloadbalancing:Describe*",
        "autoscaling:Describe*",
        "iam:GetUser",
        "dynamodb:Describe*",
        "s3:ListAllMyBuckets"
      ],
```

```
        "Resource": "*"
      }
    ]
}
```

Azure Role Setup

To allow the importing of Azure objects into the CMDB the following Azure Role Setup must be performed. Note: Azure Virtual Machine operating system discovery is performed via the Discovery App per the Windows/Linux/Unix methods listed in the Credential Requirements Table above.

1. Go to https://portal.azure.com/#blade/Microsoft_AAD_RegisteredApps/ApplicationsListBlade and create a new App Registration and note down the below details from Overview
    a. **Client ID** (Required in Virima to add Azure Credentials)
    b. **Tenant ID** (Required in Virima to add Azure Credentials)
    c. **Secret Key-Value** (Needs to be created from Certificates & Secrets) (Required in Virima to add Azure Credentials)

2. Go to https://portal.azure.com/#blade/Microsoft_Azure_Billing/SubscriptionsBlade
    a. Note down the **Subscription ID** (Required in Virima to add Azure Credentials)
    b. Go inside subscription details and click on **Access Control (IAM)**
    c. Click **Add** and then select **Add Role Assignment**
    d. Select **Reader** Role and click next
    c. Click on **'Select Members'** and search for the newly added App Registration and select it and click Select.
    e. Review and Assign the role.

User Role Types for Credentials

User Role Types	Host Scans	Pattern Scans	IP Connections	Process Scans
Standard user (No administrator access)	No	No	No	No
Local Administrator user with no elevated privileges	Yes	Yes	No	Yes
Local Administrator user with elevated privileges	Yes	Yes	Yes	Yes

System Administrator user	Yes	Yes	Yes	Yes
System Domain Administrator user (Only applicable to domain machines)	Yes	Yes	Yes	Yes

Notes:

1. Enabled Domain File sharing profile is a mandatory requirement.

2. Ivanti Neurons for Service Mapping requires admin share folder access to perform even a basic scan. This access is default with elevated privileges. Users without elevated privileges can be manually given admin share privileges by adding an entry into the registry.

3. The Discovery App need not be running on a domain machine, but if it is, the service logon credentials are required to scan that domain machine.

4. The above applies to both Windows domain and workgroup machines.

5. Credentials are stored locally within the on premise Discovery App and never shared with Ivanti Neurons for Service Mapping web services. They are encrypted and not viewable to anyone.

Additional Considerations

This document is intended to help ensure the proper scoping and sizing of Ivanti Neurons for Service Mapping as well as prepare Ivanti Professional Services and customers for initial deployment. The information provided covers the majority of customer environments and will help identify any possible exceptions that should be addressed early in the sales or professional services engagement.

Learn More

ivanti.com

1 801 308 8047 (Americas Support)

support@ivanti.com