Section 2 – How ISI-Snapshot Operates

2.1 Technical Architecture

ISI-Snapshot is a native Microsoft Windows application architecture, designed and created with Microsoft's .NET standards, components, and tools. ISI-Snapshot runs under Windows XP Professional (recommended), Windows 2000 Professional, Windows Server 2003 or later with .NET Framework 2.0.

The ISI-Snapshot collector PC communicates with the objects in the technology infrastructure to be inventoried via an IP network connection through an Ethernet connection or a Virtual Private Network (VPN) connection.

A highly normalized Microsoft SQL Server 2005 database stores the data collected during the Discovery, Credentialed Inventory, and Utilization processes. The SQL Server instance may reside locally on the ISI-Snapshot collector or on a remote server.

ISI-Snapshot utilizes different methods to communicate with components of the technology infrastructure depending on the function to be performed and the category of technology object to be inventoried.

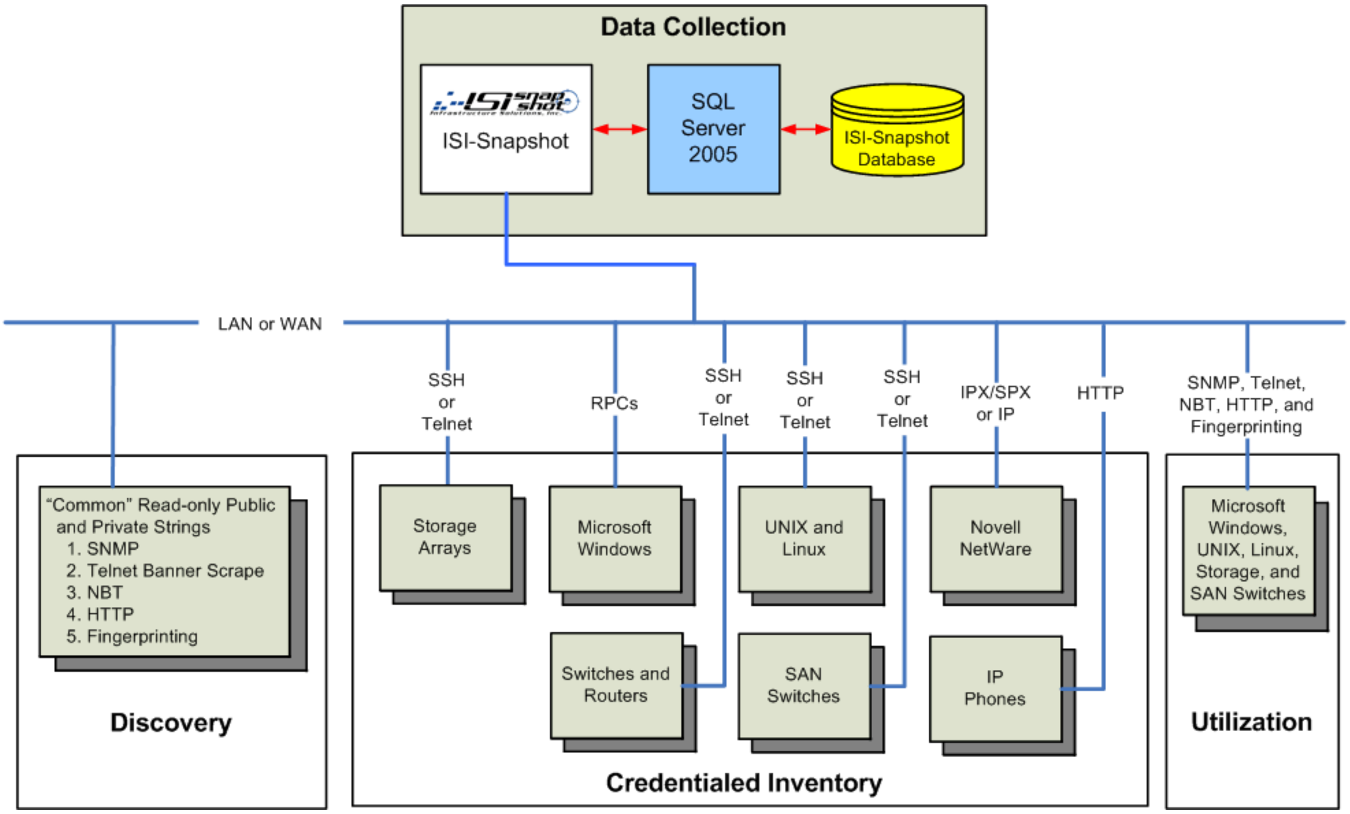

Figure 2-1. Technical Architecture Overview

2.2 Multi-threading and Dynamic Dispatching

Multi-threading enables ISI-Snapshot to perform rapid program execution by dividing the application into several independent execution paths (threads) and dispatchers that run simultaneously. Dispatchers and threads are dynamically launched, reallocated, and terminated, and how many run at any given time.

Threads are dynamically reallocated as appropriate based on real-time performance. The operator can dynamically adjust and reallocate the number of threads based on observation of performance. Multi-threading insures that dispatcher and worker threads have independent communication to middle tier components and that worker threads are completely isolated from the user interface, ISI-Snapshot database, ad-hoc query/extract facility (SnapReports), resulting in improved operator control, performance, error handling, and reliability.

eXtensible Markup Language (XML) documents are used for communication between the user interface, dispatchers, and worker threads. The use of XML documents facilitates and enables bulk database reads and writes to minimize the number of times the database is accessed and eliminates the possibility of partial or incomplete database transactions. Database writes are reduced to one per host/storage array/switch during data collection (Scan).

ISI-Snapshot implements a 3-tier design model. The three tiers are organized into:

1. User interfaces.
2. Dispatching.
3. Worker threads to perform data collection and database access.

❖ **Note:** Refer to Figure 2-2, Dynamic Dispatching and Multi-Threading, on page 2-3.

The multi-threaded design allows for significant concurrent processing limited only by the resources on the ISI-Snapshot data collection Host and the network bandwidth available. You can view the real-time performance data during collection from ISI-Snapshot and dynamically adjust, control, and optimize the level of concurrent processing and related resource consumption.

Tier 2 is the single point of communication between ISI-Snapshot and SnapReports user interfaces (Tier 1), dispatchers, worker threads (Tier 3), and database access. This tier handles and controls all database communication, validates database responses, communicates them to requesting threads and modules, and notifies the operator of exceptions.

Advantages of this design include the support of multiple concurrent users each with their own sets of user interfaces, dispatchers, and worker threads. This design also provides the capability to preview data with the query and extract facility (SnapReports) simultaneous with ongoing data collection (Scan).

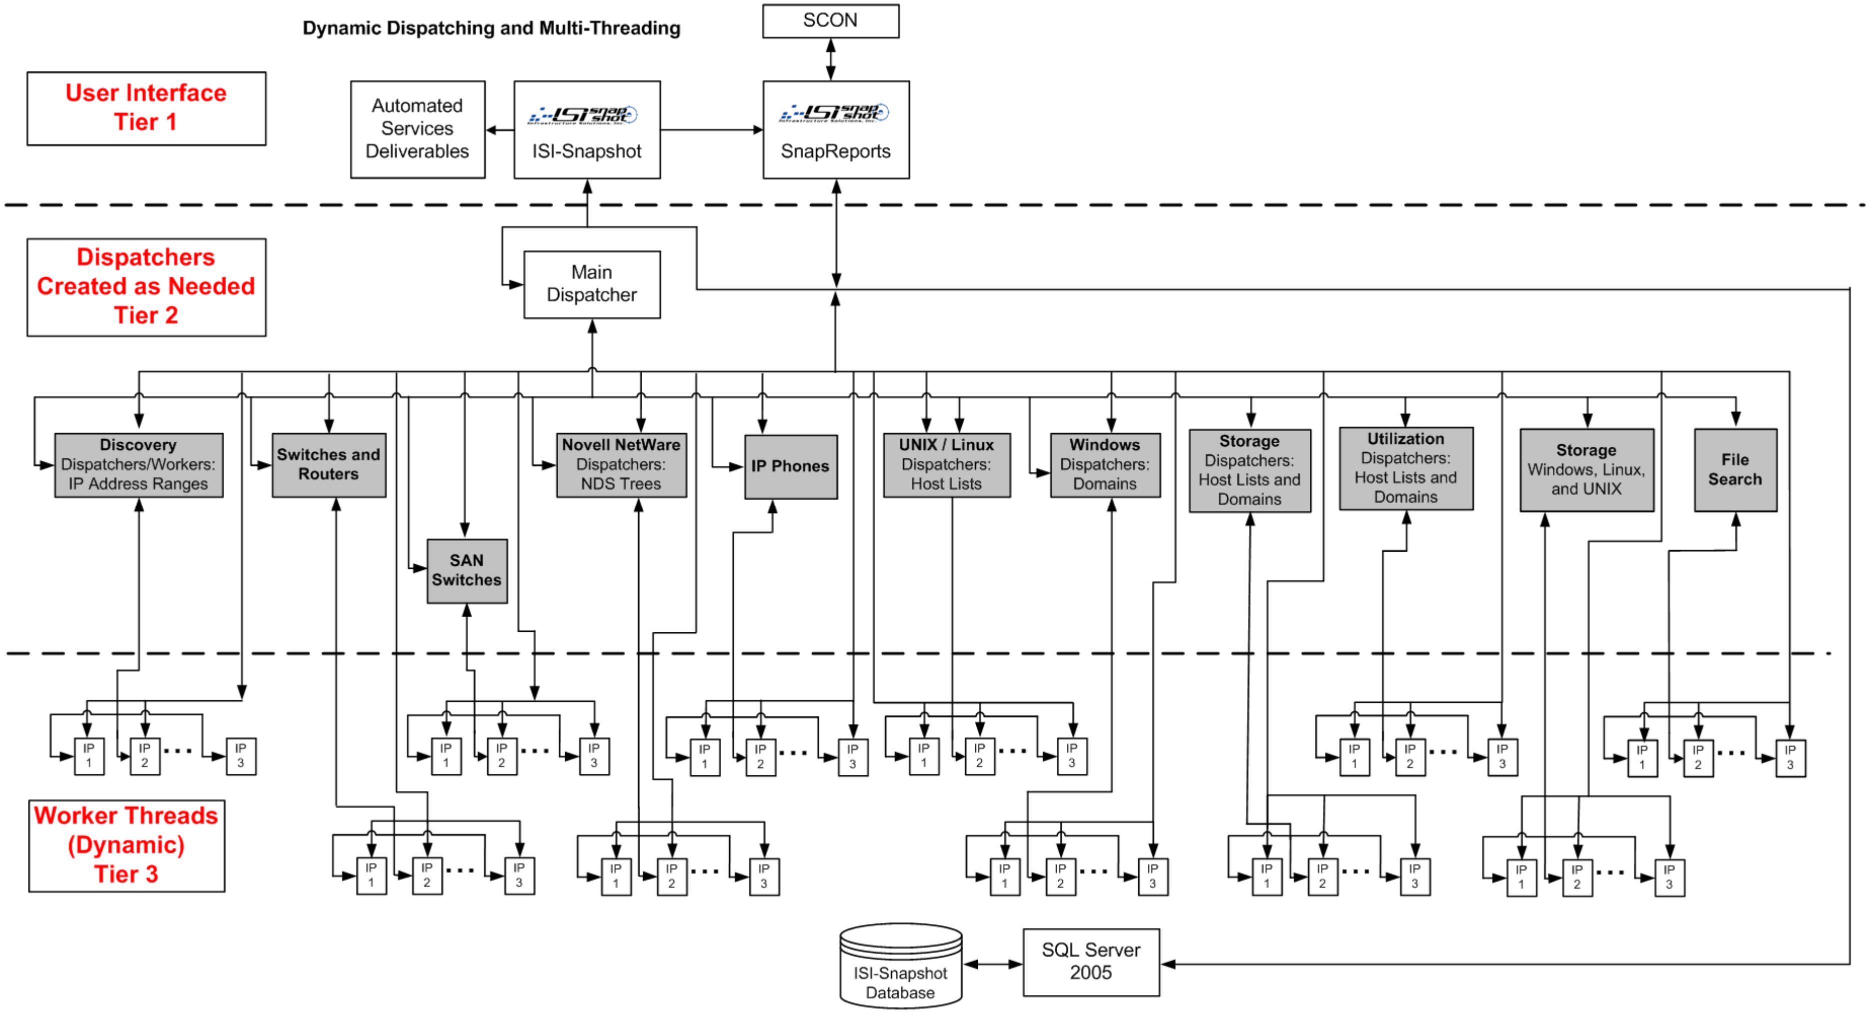

Figure 2-2. Dynamic Dispatching and Multi-Threading

2.3 Discovery (Non-Credentialed Inventory)

The ISI-Snapshot Discovery module includes functions that provide the capability to scan Internet Protocol (IP) Address ranges and attempts to discover infrastructure objects. A built-in IP Subnet/CIDR Calculator enables you to acquire IP address calculations including netmask, broadcast and network addresses including Cisco wildcard masks.

Network connectivity required by the collection host includes: unobstructed view of network, no ICMP, HTTP, RPC, or Telnet filtering.

There are multiple “levels” of connectivity that can exist between the collection host and a target network or host.

First, ISI-Snapshot must be able to route packets to the target network/device. If ISI-Snapshot cannot route to the target, then ISI-Snapshot cannot discover it. Assuming ISI-Snapshot can route packets and ping, then ISI-Snapshot can discover the device.

Next, ISI-Snapshot relies on a variety of protocols to determine information about the device. If any of these protocols are blocked (obstructed) then ISI-Snapshot risks missing an opportunity to acquire information.

The consequences of having an obstructed view of the targets are that ISI-Snapshot may not be able to detect the device exists at a particular IP address and even if ISI-Snapshot detects that the device exists, ISI-Snapshot may not be able to determine what type of device it is.

The order in which ISI-Snapshot acquires Discovery data include:

1. Simple Network Management Protocol (SNMP).
2. TELetype NETwork (Telnet).
3. NetBIOS over TCP/IP (NBT).
4. Hypertext Transfer Protocol (HTTP).
5. Fingerprinting (Port Detection).

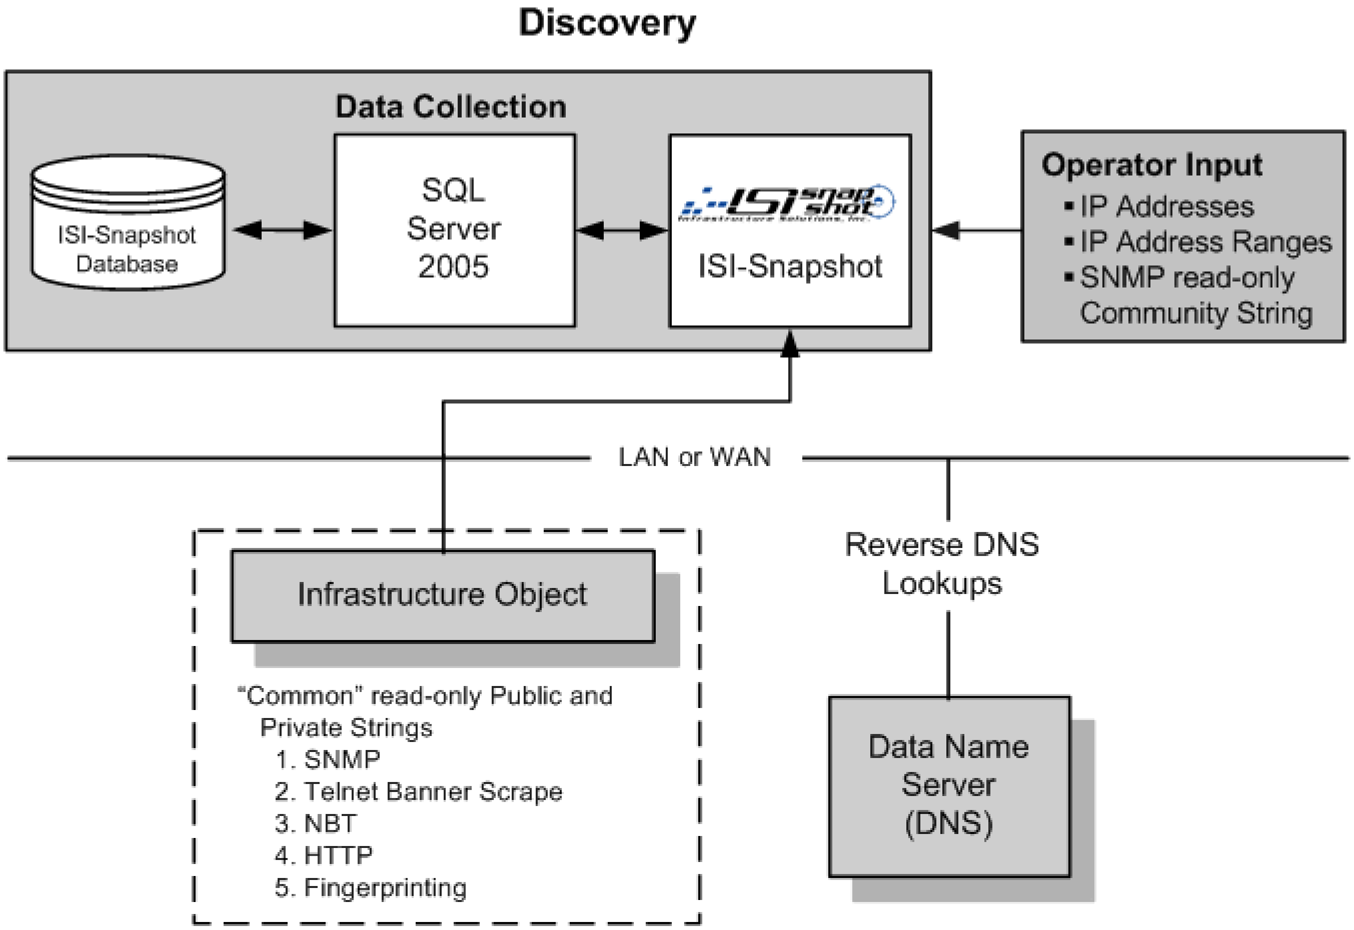

Figure 2-3. Discovery

2.3.1 SNMP (Simple Network Management Protocol)

ISI-Snapshot optionally utilizes Simple Network Management Protocol (SNMP) to rapidly interrogate and gather high level information about Internet Protocol (IP) devices attached to IP networks. You can configure ISI-Snapshot to interrogate one or more ranges of IP Addresses for active devices with SNMP agents. ISI-Snapshot uses the read-only community string that is set to “public” by device manufacturers. You can change the SNMP community strings at any time and may assign different community strings to different devices. . SNMP processing is accomplished by issuing snmpget requests.

IP devices are discovered and identified via SNMP, based on interrogation of IP Address ranges by sending an snmpget request to every user-specified IP Address. If the ISI-Snapshot collector is running an SNMP agent, ISI-Snapshot has been configured with the correct read-only community string, and the device has not been configured to respond only to specific IP Addresses, then the device responds with SNMP Management Information Base (MIB) information. The SNMP MIB-I data routinely identifies the server hardware manufacturer and model, including the operating system version and release.

You provide one or more IP Address ranges for processing. By default, IP Addresses will be skipped if they are known to exist from a previous Credentialed Inventory process.

2.3.1.1 SNMP Inventory

An snmpget request is sent to each device IP Address in the IP Address range(s). A device responds with SNMP MIB-I data only if all of the following criteria are satisfied:

- Active infrastructure object exists at the device IP Address.
- Infrastructure object has an active SNMP agent installed and running.
- Infrastructure object was not configured to respond only to clients with a limited set of IP Addresses.
- ISI-Snapshot was configured with a valid SNMP read-only community string.

2.3.1.2 SNMP Inventory Data Elements

The following MIB-I data elements are collected:

- ❖ **Note:** Values are dependent upon the System Administrator maintaining them.
- **sysDescr** – Textual description of the infrastructure object set by the manufacturer. Value is intended to include the full name and version identification of the object’s hardware type, software operating system, and networking software.
  - ❖ **Note:** This data element value rarely includes all the intended information and values lack consistency between manufacturers and even between product lines and products from the same manufacturer.
- **sysObjectID** – Vendor’s authoritative identification of the network management subsystem contained in the object. The seventh ordinal in the ObjectID uniquely identifies the device vendor values.
- **sysUpTime –** Time (in hundredths of a second) since the agent was last re-initialized.
- **sysContact –** Textual identification of the contact person for this managed node, together with information on how to contact this person.
- **sysName** – Administratively-assigned name for this infrastructure object. By convention, this is the object’s fully qualified Domain name.
- **sysLocation –** Physical location of this infrastructure object.
- **sysServices –** Value that indicates the set of services that this object may potentially offer and consists of an 8-bit value with the following bit definitions:

0	0	0	a	b	c	d	e

a: Application (NFS Servers)

b: End-to-end (IP Hosts)

c: Internal layer (IP Routers)

d: Data-link layer (Bridges)

e: Physical layer (Repeaters)

2.3.1.3 SNMP Security Requirements

SNMP agents typically support two levels of “challenge passwords,” called community strings, to authorize communication. The two types of community strings are read-only and read-write (which includes command-and-control capabilities). **ISI-Snapshot utilizes only the read-only community string**. By convention, manufacturers set the read-only community string to the value “public.” You can change the read-only community string for each device (SNMP Agent) or group of devices.

2.3.2 Telnet (TELetype NETwork)

ISI-Snapshot uses the Telnet (TELetype NETwork) protocol on Local and Wide Area Network connections. Telnet is commonly used to connect to UNIX-like operating systems, routers, switches, and other network infrastructure devices. Telnet is commonly run on TCP port 23 though it can be run on other custom ports. ISI-Snapshot Discovery attempts to establish a Telnet session and when successful, leverages the banner presented at connection time to obtain information about the device.

Because Telnet lacks any sort of transmission encryption and is susceptible to packet interception, it is used to varying degrees in the infrastructure. Lack of transmission encryption exposes the passwords and other data in the payload that allows theft by anyone who can intercept them on the network. Telnet has been superseded in many environments by SSH and other secure and encrypted transport protocols. ISI-Snapshot Inventory supports Telnet as well as SSH.

2.3.3 NetBIOS over TCP/IP (NBT)

Network Basic Input/Output System (NetBIOS) over TCP/IP (NBT or NetBT) allows legacy computer applications relying on the NetBIOS API to be used on TCP/IP networks. NetBIOS is used in Ethernet and Token Ring networks.

NBT provides three distinct services:

- Name service for name registration and resolution.
- Session service for connection-oriented communication.
- Datagram distribution service for connectionless communication.

2.3.3.1 Name Service

In NBT, each participant must register on the network using a unique name of at most 16 characters. NBT implements a central repository (Name Service) that records all name registrations. An application wanting to register a name contacts the name server (which has a known network address) and asks whether the name is already registered, using a “Name Query” packet. The name server returns a negative response immediately if the name is not already in the database, meaning it is available.

The packet formats of the Name Service are identical to DNS. The key differences are the addition of NetBIOS “Node Status” query, dynamic registration, and conflict marking packets. They are encapsulated in User Datagram Protocol (UDP).

In addition, to start a session or to send a datagram to a particular Host rather than to broadcast the datagram, NBT determines the IP Address of the Host with a given NetBIOS name using the Name Query packet. The response displays the IP address of the Host with that name.

2.3.3.2 Session Service

Session mode lets two computers establish a connection for a “conversation,” allows larger messages to be handled, and provides error detection and recovery. Sessions are established by exchanging packets. The computer establishing the session attempts to make a TCP connection to port 139 on the computer with which the session is to be established.

If the connection is made, the computer establishing the session then sends over the connection a “Session Request” packet with the NetBIOS names of the application establishing the session and the NetBIOS name to which the session is to be established. The computer with which the session is to be established responds with a “Negative Session Response” indicating that no session can be established (either because that computer is not listening for sessions being established to that name or because no resources are available to establish a session to that name) or a “Positive Session Response” indicates that a session can be established.

Data is transmitted during an established session by Session Message packets. TCP handles flow control and retransmission of all session service packets, and the dividing of the data stream over which the packets are transmitted into IP datagrams small enough to fit in link-layer packets. Closing the TCP connection closes sessions.

2.3.3.3 Datagram Distribution Service

Datagram mode is “connectionless” (each message is sent independently), messages must be smaller, and the application is responsible for error detection and recovery. NetBIOS datagrams are sent over User Datagram Protocol (UDP). A datagram is sent with a “Direct Unique” or “Direct Group” packet if it is being sent to a particular NetBIOS name, or a “Broadcast” packet if it is being sent to all NetBIOS names on the network.

2.3.4 Hypertext Transfer Protocol (HTTP)

ISI-Snapshot uses the Hypertext Transfer Protocol (HTTP) communications to acquire data from devices across Internet services. HTTP is a request/response protocol between clients and servers. The originating client, such as a web browser, spider, or other end-user tool, is referred to as the user agent. The destination server, which stores or creates resources, such as HTML files and images, is called the origin server. In between the user agent and origin server there may be several intermediaries, such as proxies, gateways, and tunnels. HTTP does not require TCP/IP. HTTP can be implemented on top of any other protocol on the Internet, or on other networks. HTTP only requires a reliable transport so you can use any protocol that provides such.

An HTTP client initiates a request by establishing a Transmission Control Protocol (TCP) connection to a particular port on a remote device (port 80 by default). An HTTP server listening on that port waits for the client to send a request message.

Upon receiving the request, the server sends back a status line, such as “HTTP/1.1 200 OK”, and a message of its own, the body of which is the requested file, an error message, or some other information.

2.3.5 Fingerprinting

If the previous Discovery methods fail, ISI-Snapshot leverages an assortment of methods including port scans and “fingerprinting technology” to gather as much information as possible about the device. Fingerprinting interrogates one or more IP Address ranges. During Discovery, ICMP echo and TCP Ping requests are sent to each device Address in the IP Address range. If the device returns a positive acknowledgement, a series of port scans attempts for various device ports. Fingerprinting requires name resolution configuration for the ISI-Snapshot collector PC.

2.4 Credentialed Inventory

Credentialed Inventory is the capability to verify that an infrastructure object with a specific IP Address exists and to successfully present authorized administrative security credentials, such as user name and password. This allows administrative access to the object in order to capture and retrieve configuration and utilization data without adversely effecting its operation and state.

In the case of servers, Credentialed Inventory consists of successfully logging in to an administrative account on the device and executing a series of read-only or display commands, Application Program Interface (API) requests, or Remote Procedure Calls (RPCs) to retrieve configuration and selected utilization data.

Credentialed Inventory includes Host, Switches and Routers, SAN Switches, IP Phones, and Storage Array inventories.

2.4.1 Host Operating Systems

Host Operating Systems supported by ISI-Snapshot Credentialed Inventory include:

- Microsoft Windows Servers and Workstations.
- UNIX and Linux Systems.
- Novell NetWare Servers.

Host credential inventories collect data detailing the physical and logical assets of a server.

2.4.2 Storage Arrays

ISI-Snapshot collects configuration data from SAN-attached “enterprise” storage arrays, and Network Attached Storage (NAS) devices. This type of Inventory collects data detailing disk drives that appear to be Host systems and various storage peripherals. The actual file system is accessed for Inventory.

2.4.3 Switches and Routers

ISI-Snapshot collects configuration data from Cisco Switches and Routers using the same Credentialed Inventory methodology used for the collection of UNIX and Storage devices. ISI-Snapshot connects to the Cisco devices using the Telnet or SSH protocol. ISI-Snapshot requires a username and/or password to login to Cisco devices and execute IOS Cisco “show” command. At the start of the collection Scan, an ICMP echo is sent to the Cisco device; once connected to the Cisco device, ISI-Snapshot runs a series of “show” commands and populates the database with the various responses received from the Cisco device.

As with the other supported platforms, flexible configuration settings allow for the changing of timeout values, prompts, ICMP Timeouts and Attempts as well as an extended debugging option to aid in the data collection process. Options are available to collect additional information on the Cisco devices including CDP Neighbors, Switch and Router Interfaces, Line Cards, and Routing Tables. By selecting a checkbox in the Cisco Scan Settings, ISI-Snapshot collects the additional information based on the selection. This information is available for reporting along with the base information collected from the Cisco devices and can be seen in different Views and Reports.

2.4.4 Storage Area Network (SAN) Switches

ISI-Snapshot inventories Storage Area Network (SAN) Switches that attach remote computer storage devices, such as disk array controllers, tape libraries and CD arrays to servers in such a way that to the operating system the devices appear as locally attached devices.

By contrast to a SAN, Network-Attached Storage (NAS) uses file-based protocols, such as NFS or SMB/CIFS where it is clear that the storage is remote, and computers request a portion of an abstract file rather than a disk block. ISI-Snapshot supports the collection of configuration data from Brocade, Cisco MDS 9000, EMC Connectix, and McDATA SAN Switches and Directors.

2.4.5 IP Phones

ISI-Snapshot collects configuration data from Cisco 7900 Series IP Phones. All that is required for a collection of Cisco IP Phones is a range of IP addresses that ISI-Snapshot can process during the Scan. ISI-Snapshot attempts to connect to each of the IP Addresses and collects basic configuration data from each Cisco IP Phone. A View has been added to the Reports module that displays the information collected from the Cisco IP Phones.

2.4.6 Microsoft Windows Servers and Workstations

Inventory of Windows servers and workstations is accomplished by merging data from the Windows Browse List, Windows Active Directory, any Windows NT Security Accounts Manager (SAMs), and any operator supplied Host Lists and exclusions. Once the set of Windows devices is obtained, ISI-Snapshot communicates with Microsoft Windows devices via Remote Procedure Calls (RPCs) to log in and capture configuration and Inventory data. Windows authentication sends the credentials of the user logged into the collection host before sending configured credentials.

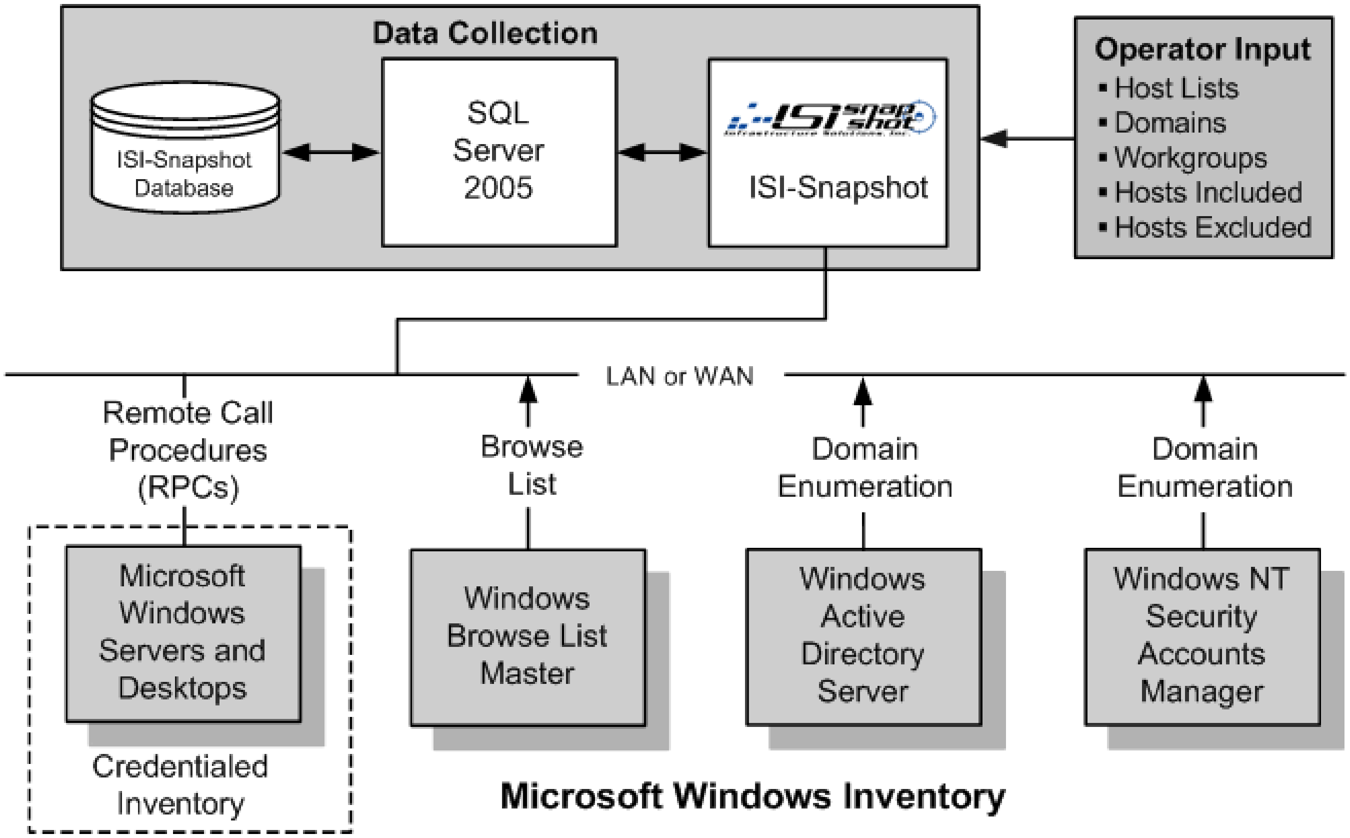

Figure 2-4. Microsoft Windows Inventory

2.4.6.1 Windows Host Enumeration Method

Windows devices are inventoried by acquiring and merging data from four potential sources of Host information depending on user specifications:

- Windows Browse List.
- Enumerating Windows Active Directory Domain information.
- Enumerating (any) Windows NT Security Accounts Manager (SAM) databases.
- Processing (any) selected Windows Host Lists.
- Host List from an ISI-Snapshot Discovery report from a prior Run.

You control the process by selecting one or more Domains, Workgroups, and Hosts for processing. Excluding specific Domains, Workgroups, and Hosts can further control the process.

2.4.6.2 Windows Inventory

ISI-Snapshot attempts to log in to each Windows Host to be inventoried with the administrative security credentials. If log in is successful, a series of Windows APIs executes to retrieve configuration and select utilization data. Communication with target device is accomplished via Remote Procedure Calls (RPCs). Windows authentication sends the credentials of the user logged into the collection host before sending configured credentials.

2.4.6.3 Windows Security Requirements

ISI-Snapshot requires a log in to an administrative account on each Windows Host to be inventoried. The account must have administrator authority for ISI-Snapshot to function properly.

Servers in Domains – For Windows devices in Domains, the log in account should have Domain administrator authority. In Windows infrastructures with more than one Domain, the Domains may be implemented with full (two-way) trust relationships. All servers in trusting Domains that have the "Domain Administrator" group from the trusted Domain in the "Local Administrator" group on servers in the trusting Domain can be successfully inventoried with a single administrative security account. For servers in other Domains, separate administrative security accounts must be provided for each Domain.

- ❖ **Note:** If your Domain administrator accounts have been disabled or removed from servers in Domains, the only way to successfully log in and Inventory the servers is to log in to each individual Host using an account with local administrator security credentials.

Servers in Workgroups – For Windows devices in Workgroups, ISI-Snapshot must log in to each individual Host using an account with local administrator security credentials. Depending on the policies, standards, conventions, and practices in use, this may require a unique local administrator account on each Host to be inventoried.

The following services and facilities must be running and enabled:

- ♦ Remote Procedure Call (RPC) Service must be running and available to ISI-Snapshot using the administrator account ISI-Snapshot has been given to use for Inventory.
- ♦ Remote Registry Service must be running.
- ♦ ISI-Snapshot must have the entire HKEY_LOCAL_MACHINE hive readable on the target from the computer running ISI-Snapshot. Certain local/domain Windows policy restrictions can prevent this access, which causes the Inventory to fail even though Remote Registry is running and available.
- ♦ Server Service must be running.
- ♦ File and Print Sharing for Microsoft Networks must be enabled.
- ♦ Default administrative shares ADMIN$ and IPC$ must be defined and enabled.

2.4.6.4 Windows Network Port Requirements

Services on Targets:

- Computer Browser.
- Server Service.
- Remote Registry Service.
- WMI (Windows Management Instrumentation).
- Task Scheduler Service.

Network Requirements on Targets:

- TCP/IP Protocol.
- File and Printer Sharing for Microsoft Networks.
- NetBIOS over TCP/IP enabled.
- Target must also be responsive to an ICMP request (ping).

2.4.7 UNIX and Linux Systems

Inventory of UNIX and Linux servers is accomplished primarily by relying on operator supplied Host Lists and exclusions. Once the set of Hosts is obtained, ISI-Snapshot communicates with UNIX and Linux Hosts via SSH or Telnet to log in and capture configuration and Inventory data by executing read-only display and other commands.

The Network Information Service (NIS) is UNIX and Linux client-server directory service protocol for distributing system configuration data, such as user and host names between computers on a computer network. ISI-Snapshot uses NIS or NIS+ to Inventory all systems for use in the staging method.

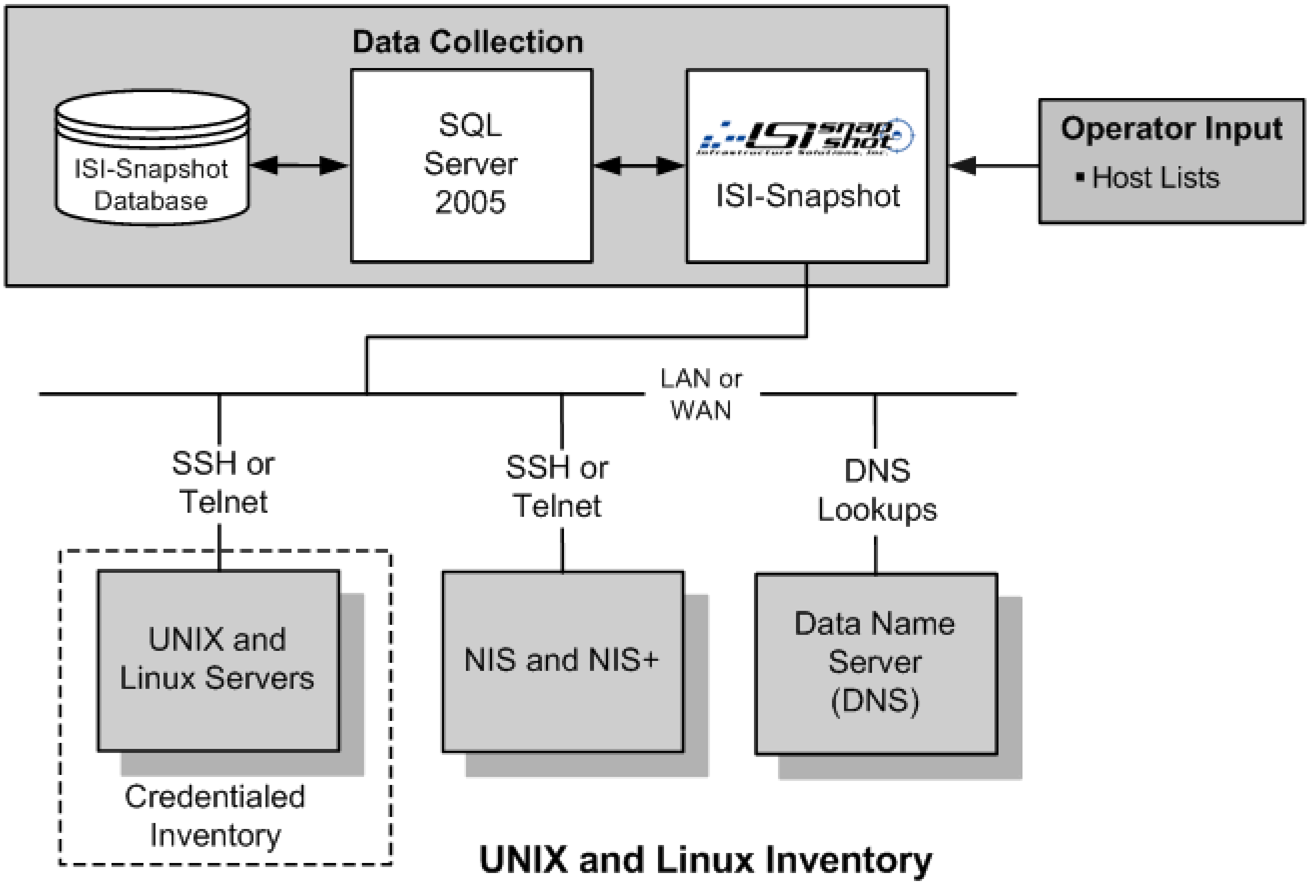

Figure 2-5. UNIX and Linux Inventory

2.4.7.1 UNIX and Linux Host Enumeration Method

UNIX and Linux servers to be inventoried are identified by one or more lists of servers and the process works best by grouping servers that share common security approaches and credentials into the same Host List.

2.4.7.2 UNIX and Linux Inventory

ISI-Snapshot attempts to log in to each UNIX and Linux Host to be inventoried with administrative security credentials provided. If log in is successful, a series of commands executes to retrieve configuration and selected utilization data. Communication with target servers is accomplished via Secure SHell (SSH) or Telnet communications. The default IP Ports for SSH and Telnet is used unless you select different Port numbers. ISI-Snapshot can be configured to exclusively utilize either SSH or Telnet. ISI-Snapshot automatically detects which Port (SSH or Telnet) is available and uses that Port. If both Ports are available, ISI-Snapshot selects SSH.

2.4.7.3 UNIX and Linux Security Requirements

ISI-Snapshot requires log in credential to an account on each UNIX and Linux Host to be inventoried. The account must have root authority for ISI-Snapshot to function properly. In cases where policies prevent direct log in as root, ISI-Snapshot provides the capability to log in to a user account and switch to root with the SU (Substitute User or Super User), sudo, PowerBroker, or any custom commands.

ISI-Snapshot provides support for the following UNIX and Linux security models:

- **Direct log in to root account on servers –** ISI-Snapshot can log in via SSH or Telnet to the root account on a UNIX or Linux Host.
- **Log in to user account and SU to root account –** ISI-Snapshot can log in via SSH or Telnet to a user account on a UNIX or Linux Host. Then the SU command can be executed to switch to the root account.
- **Log in to user account and sudo to root account –** ISI-Snapshot can log in via SSH or Telnet to a user account on a UNIX or Linux Host. Then the sudo command can be executed to switch to the root account.
- **Trusted Host support –** ISI-Snapshot provides support to log in to UNIX and Linux servers via their Trusted Hosts.
- **SSH with encrypted keyfile(s) –** ISI-Snapshot provides support to utilize SSH encrypted keyfiles to gain administrative access to UNIX and Linux systems.
- **Lightweight Directory Access Protocol (LDAP) –** ISI-Snapshot is compatible with LDAP security implementations that may be in use within the UNIX and Linux technology infrastructure.
- **NIS and NIS+ –** ISI-Snapshot is compatible with Network Information Service (NIS) and NIS+ security implementations that may be in use within the UNIX and Linux technology infrastructure.
- **PowerBroker –** ISI-Snapshot is compatible with PowerBroker that securely delegates root and other special account privileges.
- **User Define Authorization Command –** ISI-Snapshot provides support for the CMD: command and executes the command entered after the colon.

2.4.7.4 UNIX and Linux Network Port Requirements

Services/Daemons on Targets:

- Telnet Service.
- SSH Service.

Network Requirements on Targets:

- TCP/IP Protocol.
- Target must also be responsive to an ICMP request (ping).

2.4.8 Novell NetWare Servers

Inventory of Novell NetWare servers merges information from Novell Directory Services (NDS) with user supplied Host Lists and exclusions. NDS is required. Once the set of servers is obtained, ISI-Snapshot communicates with Novell NetWare servers via IPX/SPX or IP to log in and capture configuration and Inventory data by executing read-only display and other commands.

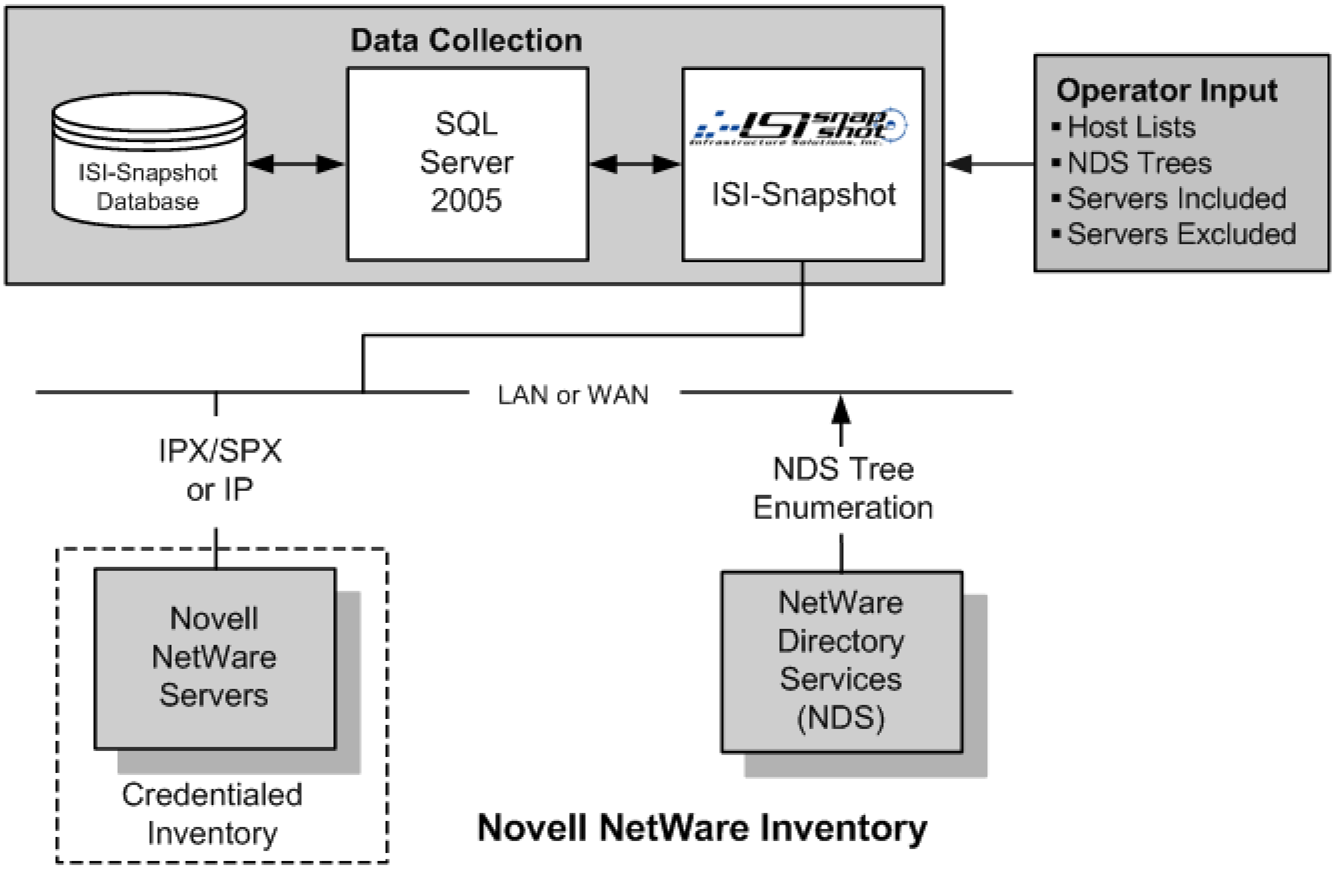

Figure 2-6. Novell NetWare Inventory

2.4.8.1 NetWare Host Enumeration Method

NetWare servers are discovered by gathering and merging data from two potential sources of Host information depending on operator specification:

- Novell Directory Services (NDS) Trees.
- Processing (any) selected NetWare Server List(s).
- Host list.

Selecting one or more NDS Trees for processing controls the Discovery process. Excluding specific NDS Trees and servers can further control the process. If IP is the only protocol in use, you must specify the “scope” and Service Location Protocol (SLP) directory agent addresses to enumerate NDS Trees and servers.

2.4.8.2 NetWare Inventory

ISI-Snapshot attempts to log in to each NetWare Host to be inventoried with the administrative security credentials provided. If log in is successful, a series of NetWare APIs executes to retrieve configuration and selected utilization data. Communication with target servers is accomplished via Novell’s NetWare Client32 software, which utilizes IPX/SPX or IP, depending on the NetWare configuration.

2.4.8.3 NetWare Security Requirements

ISI-Snapshot requires access to log in to an account on each NetWare server to be inventoried. The account must have administrator authority for ISI-Snapshot to function properly. This means the account must be an “admin equivalent” account, which is a member of the NetWare “console operator” group and a “trustee” of the NDS Tree root.

2.4.8.4 NetWare Network Port Requirements

- 213 IPX Protocol Credentialed Inventory of Novel NetWare using IPX/SPX.
- 524 NCP NetWare Core Protocol Credentialed Inventory of Novel NetWare using TCP/IP.

2.5 Utilization for Collection

The ISI-Snapshot Utilization feature performs limited performance monitoring to collect resource Utilization data that is stored in the ISI-Snapshot database and displays in the SnapReports grid and the Utilization Manager.

Utilization polls the devices and returns data including processes running, percentage of disk busy, IP network Utilization, disk inputs/outputs (I/O), and counters. Utilization executes on the most recently closed Run.

Counters are categories and instances on which to scan. Counts are returned in the appropriate descriptor for the counter selected.

❖ **Note:** Utilization counters display the continuous performance counters over a period of time and NOT at a particular point in time.

ISI-Snapshot Utilization supports the collection of Utilization data from Windows servers and workstations, Linux systems, UNIX systems, Brocade, Cisco MDS 9000, EMC Connectix, and McDATA SAN Switches and Directors. Refer to Section 2.4.4, Storage Area Network (SAN) Switches, on page 2-10.

Utilization collection configuration provides a Scheduler that allows you to define both the start time and intervals for Utilization snapshots.

Once Inventory and Utilization have occurred, you can use the data and the Server Consolidation Application (SCON) to accurately forecast resource Utilization between different configurations. Configuration Plans can be stored based on the Consolidation Scenarios you develop.

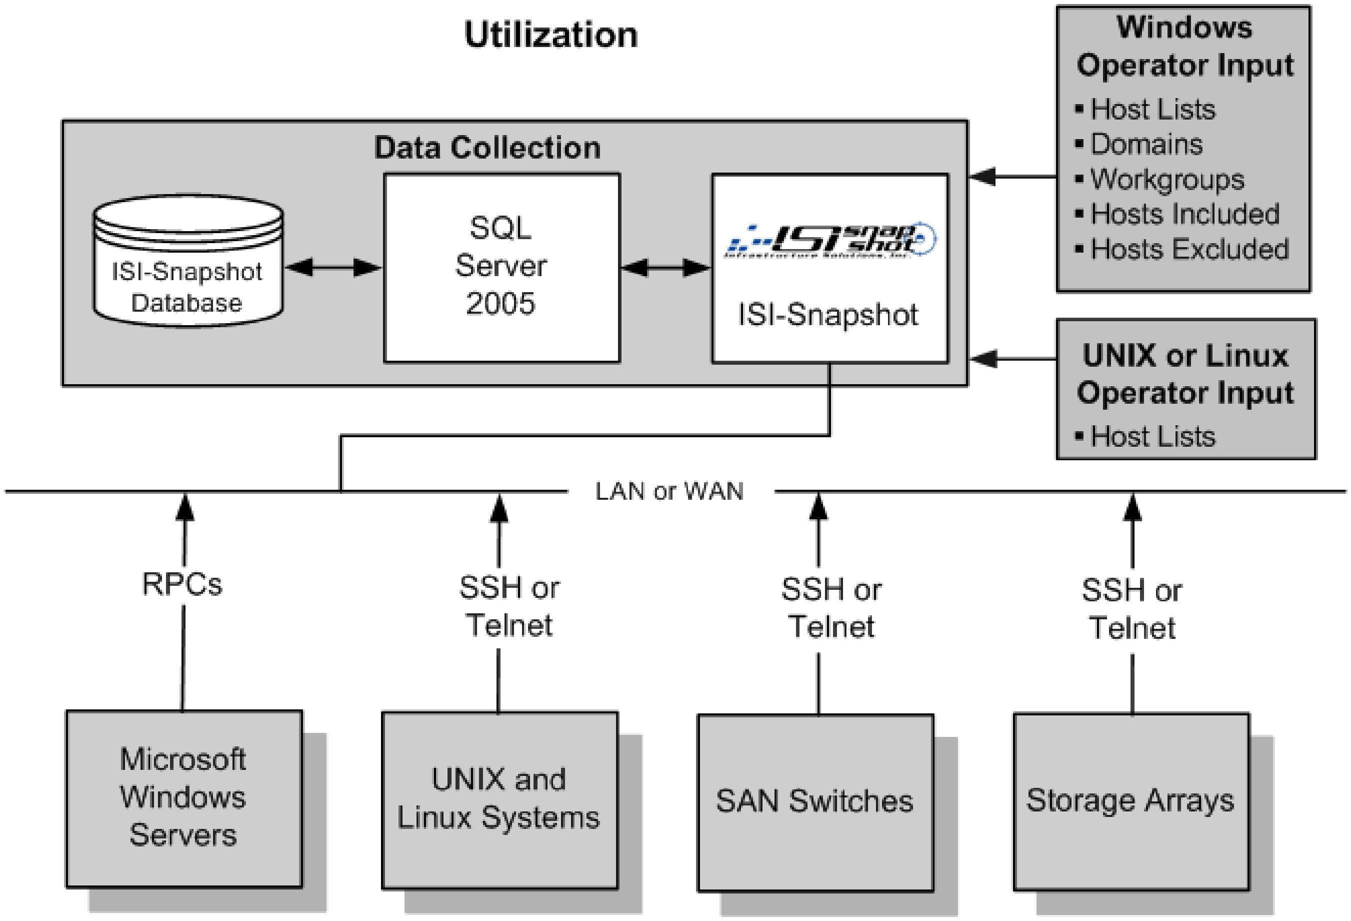

Figure 2-7. Utilization Performance Monitoring

2.5.1 Utilization Attributes Collected by Platform

Utilization hosts supported by ISI-Snapshot include:

- Microsoft Windows Servers and Workstations.
- UNIX and Linux Systems.
- Storage System Arrays.
- SAN Switches.

2.5.2 Utilization Counters

The Table 2-1, Table 2-2, and Table 2-3 provide the Utilization counters available for collection by ISI-Snapshot.

Table 2-1. Utilization for Windows, UNIX, and Linux

Counter	Type
Processor	▪ % Processor Time
Memory	▪ % Committed Bytes In Use ▪ Available Bytes ▪ Pages/Second
Paging File	▪ % Usage
Process	▪ % Processor Time Working Set
Physical Disk	▪ % Disk Time ▪ Current Disk Queue Length
Network Interface	▪ Bytes Total/Second
HBA	▪ PortStats

Table 2-2. Utilization for NetApp Storage Arrays

Counter
SYSStat
LUNStat
FCPStat
iSCSIStat

Table 2-3. Utilization for SAN Switches

Counter	Type
Processor	▪ % Processor Time
Memory	▪ % Committed Bytes In Use ▪ Available Bytes
Port	▪ Input Throughput Mbps ▪ Output Throughput Mbps ▪ OLS In ▪ OLS Out ▪ LRR In ▪ LRR Out ▪ Link Failures ▪ Synch Losses ▪ Signal Losses

2.6 Group Management

Groups and Categories can be created, edited, or deleted using Group Management. Use a Category to define what types of Groups are maintained. Categories include Business Units, Environment, Location, Primary Functions, or other user-defined Categories.

For example, the Category Business Units could include Groups called Human Resources, Finance, Corporate, or Accounting. The Accounting Business Unit may have two sub groups called Accounts Receivable and Account Payable.

You can create customized Categories and Groups. Use Groups and Categories to view forecast resource Utilization between different configurations of existing servers or new, user-defined servers.

Groups and Categories can be created, edited, or deleted using Group Management. Use a Category to define what type of Groups are maintained. Categories include Business Units, Environment, Location, Primary Functions, or other user-defined Categories.

Group Management provides the following capabilities:

- Operator defined Groups and Subgroups.
- Capability to assign servers to one or more Group and Subgroup memberships.
- Ability to use Group memberships as selection criteria for server Consolidation Candidates.

2.7 Server Consolidation Application (SCON)

Once Inventory and Utilization data has been collected, you can use the data and the Server Consolidation Application (SCON) to accurately forecast resource Utilization between different configurations called Consolidation Scenarios. Consolidation Scenarios can be saved as Consolidation Plans. SCON allows you to consolidate servers reducing costs and resources, such as man-hours, hardware, maintenance, and replication.

SCON provides cross-platform support with the ability to analyze potential consolidation across operating system versions and the same or different platforms.

SCON creates the following results and outputs:

- **Target Utilization –** As existing server workloads are consolidated to the Target Server, a set of Performance Meters are adjusted in real time to display Target Server processor, memory, and network interface utilization. This also occurs as each change is made to the target configuration.
- **Target and Source Utilization –** Peak and historical.
- **Consolidation Reporting –** Including Charting and Print Preview support.

2.8 Automated Professional Services Deliverables (APSD)

Automated Professional Services Deliverables (APSD) enables the creation of customized service deliverables documentation driven by ISI-Snapshot data, user defined business rules, and custom document templates. This information automatically produces the data driven customer deliverable documents. APSD includes:

- A platform for the fast and easy creation of professional services assessment documentation.
- Low product complexity and highly automated features allow this solution to be implemented and delivered by the Technical Sales Team.
- A rules driven solution with rule and output content owned by the partner.
- “Real data” collected in the client’s environment and the partner’s intellectual capital as reflected in the rules.
- Ability to integrate directly with ISI-Snapshot.
- The use of industry standard formats.
- Rules created in XML or Microsoft Excel and are proprietary to your organization.
- Document templates and output content uses Microsoft Office (Word, Excel, and PowerPoint).
- Documents that can be edited or modified post creation without special tools.

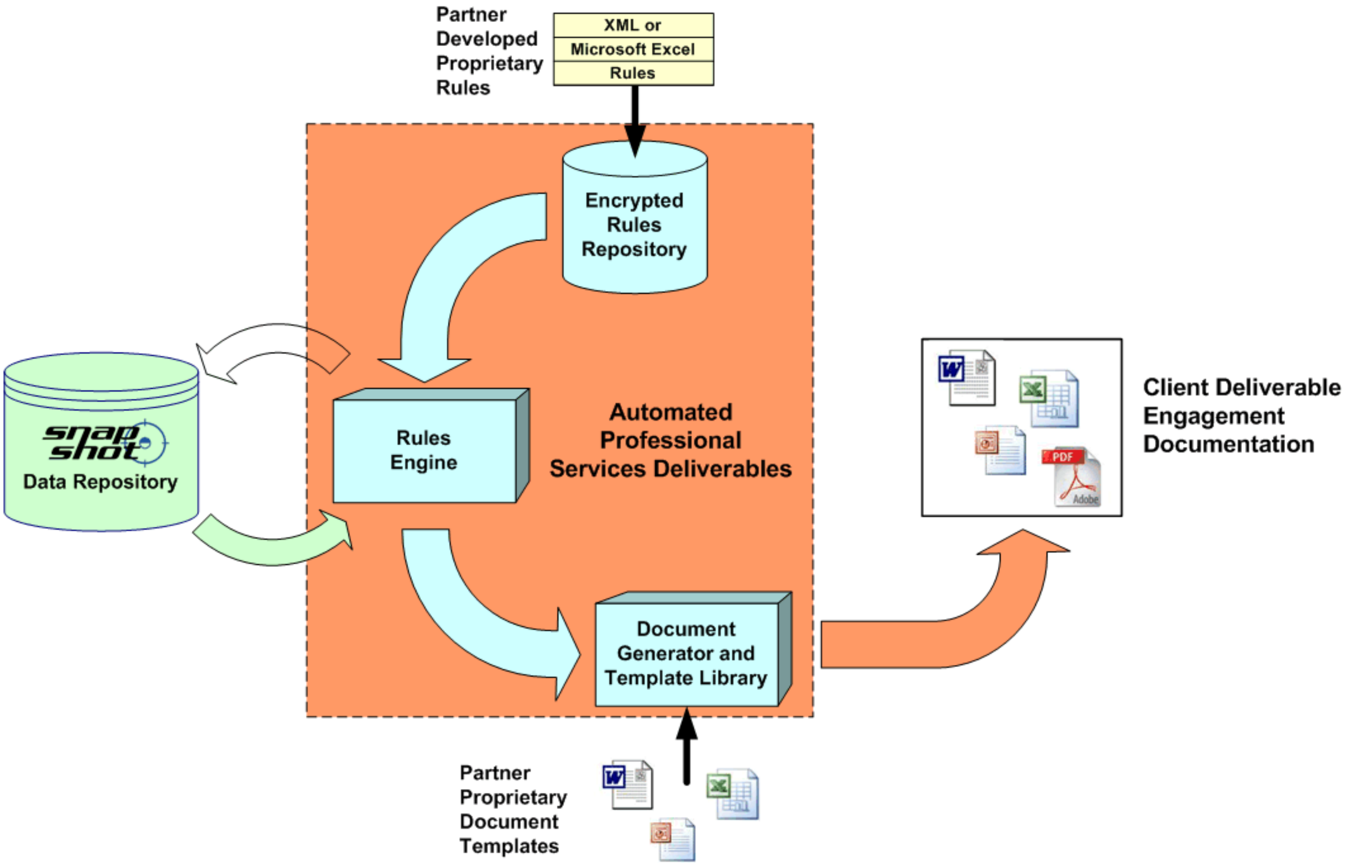

Figure 2-8. Automated Professional Services Deliverables (APSD)