Check Point
SOFTWARE TECHNOLOGIES LTD.

CLOUD • MOBILE • THREAT PREVENTION

WELCOME TO THE FUTURE OF CYBER SECURITY

CHECK POINT
Identity Awareness
Reference Architecture and Best Practices

INTRODUCTION

There is a wealth of contextual metadata available about network devices once they join a network. Traditional firewalls that enforce security policies defined with IP addresses are largely unaware of the user and device identities behind those IP addresses. They rely on static rule bases and are unable to enforce dynamic users and role-based access, or provide important metadata and context in logs and security reports.

> “New partnership and customer engagement models have extended the identity boundary of today's digital businesses: Security pros must manage identities and access across a variety of populations (employees, partners, and customers), device access methods, and hosting models. A strong digital IAM strategy protects the firm and its customers from sophisticated cybercriminals and improves efficiencies.” Forrester 2018 [1]

THE IDENTITY PROCESS

Network Authentication, Authorization, and Accounting (AAA, pronounced "triple-A") has been in use since the dawn of the Internet. Authentication asks the question, "Who or what are you?" Authorization asks, "What are you allowed to do?" And finally, Accounting wants to know, "What did you do?"

Authentication – Who or What are You?

Check Point devices usually learn who the user is from other devices, but there are cases where Check Point authenticates the user, e.g. through remote VPN access or a captive portal where the user’s network request is redirected to a browser-based authentication page. The client platform may be Windows, macOS, Linux, Android or Apple iOS. The authentication method may be one or more of a username/password or a digital certificate in a Check Point database. Check Point also authenticates users against external stores, such as RADIUS, SecurID, LDAP and TACACS (sk149854).

Authorization – What are You Allowed to Do?

What the user is allowed to do depends upon rules in our Unified Access Control Policy where the source can be an Access Role object. An Access Role is a logical representation of users and devices comprised of four elements: networks @ user or group @ machine @ client, where client is one of the Check Point remote access clients.

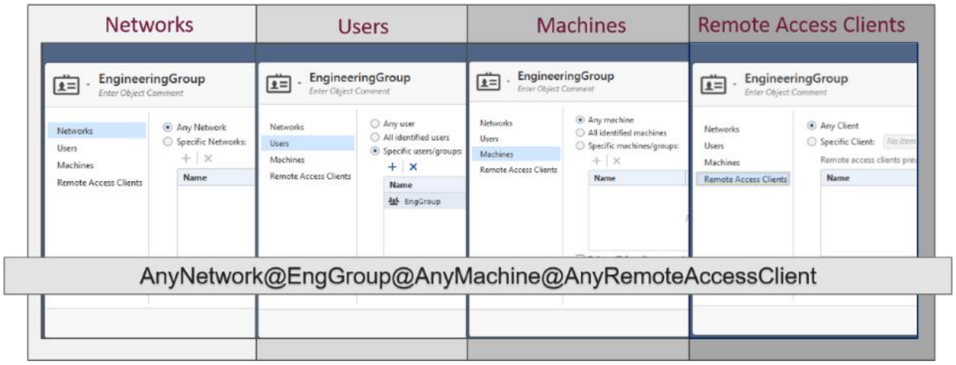

For example, it is possible to configure an Access Role object describing a User/Group object that is AND related with the IP address range of a network. An Access Role object like this will only match if both dimensions match: a user is part of a gro up AND the connection is initiated from the IP network range configured.

Further defining what the user is authorized to do, other components of a Unified Access Control Policy rule may be an IP address or network, a network port or service, an application or URL, a file or data type and the time of day. If Mobile Access is enabled, the policy may also include a mobile application; such as Web applications, file shares, Citrix services, Web mail, and native applications.

Accounting – What Did You Do?

User activity is captured in audit and security logs enhanced with contextual content from our integrated security products.

IDENTITY AWARENESS

In most environments, a Check Point security gateway will not directly handle the access request to join a network. For instance, Windows endpoints will log on to the network as part of a Windows Active Directory Domain. Mobile devices and guest users may access the network when connecting to a Wi-Fi access point. A terminal server may handle multiple connections to multiple client systems, connecting them to the network. Larger organizations may have NAC solutions in place, such as a Cisco TrustSec deployment. In these cases, Check Point can acquire identity elements from these other sources to enforce user-based policy decisions.

Identity Sources

Check Point Identity Awareness works well in these environments. We connect to these external components to map the users or devices and their IP addresses that define Access Role elements and to gain additional metadata from the source.

Directory Services	Agents	RADIUS Accounting or Identity Awareness Web API	Syslog Integrations	Integrations with Third Party APIs
Active Directory NetIQ eDirectory	Check Pont Remote Access Client Check Point Identity Agent Check Point Terminal Server Agent	Cisco Wireless LAN Controller Cisco ISE Aruba ClearPass Forescout CounterAct F5 Pulse Secure SilverFort SecurePush	Cisco ASA Fortinet	Cisco TrustSec Pulse Secure

As you can see, Check Point has several methods for connecting to various identity sources such as using RADIUS accounting and parsing syslog messages. In addition, other vendors and third parties can manage identity elements using our Check Point Identity Awareness Web API.

Excluding the Check Point clients captive portal and remote access discussed in the "Authentication – Who or What are You?" section above, these are the connectors available.

	Description
Terminal Servers	Identities are acquired using agents installed on a Windows-based application server that hosts Terminal Servers, Citrix XenApp, and Citrix Xen Desktop services. These agents identify individual user traffic coming from Terminal Servers.
Identity Agents	Identities are acquired using agents that are installed on the Endpoint computers.
Active Directory Query	Identities are acquired seamlessly from Microsoft Active Directory WMI API.
RADIUS Accounting	Identities are acquired from a RADIUS accounting client.

Identity Collector	Identities are acquired using a multi-purpose agent installed on a Windows host. The agent uses APIs to connect to Microsoft Active Directory Domain Controllers, Cisco ISE servers, and NetIQ eDirectory LDAP servers. The agent can also parse syslog messages to extract identities from a syslog message.
Identity Web API	Gives you a flexible method for creating identities and easily perform third party integrations.

Identity Sessions

These connectors provide the Identity Sources mapping metadata to the Policy Decision Point (PDP) process running on Check Point security gateways, where learned identities are stored using a session model. Sessions save the user, machine, IP address, groups and access roles for each associated identity, which will help us to determine the needed enforcement for this entity.

ⓘ The number of identities a gateway can manage is limited by kernel tables. By default, the related tables support up to 30,000 identities and can be increased to up to 200,000 in R80.10 and above.

Group Resolution

Not every identity source provides all of the elements needed to define an access role. For instance, some identity sources do not provide the group that the user is a member of. This is an important part of most Access Role definitions because group categorization simplifies the security policy, resulting in a more dynamic policy.

RADIUS Accounting	Identity Collector with Cisco ISE	IDA Web API	Captive Portal	AD Query	Identity Collector without Cisco ISE	Terminal Server Agent	Identity Agent
Can provide additional group information			A separate group query is needed				

If the group is unknown, but defined in the Access Role object, then the PDP does an LDAP query to find the group membership of the user. For instance, in an Active Directory environment when a first logon event occurs and there is no entry in cache, then an LDAP group membership query is sent to the AD server.

Identity Awareness also supports LDAP nested groups. When a group is nested in another group, the users in the nested group are identified as part of the parent group. There are three ways to configure nested group queries:

- Recursive (default): The gateway sends up to 20 queries if there are 20 or more groups that the user belongs to (sk66561).
- Per-user (from R80.10): With one LDAP query, the response includes all groups for the given user, with all nesting levels.
- Multi per-group (from R80.10): The gateway sends one LDAP query, which includes the user name and the group. The server responds if the user is or is not in this group.

ⓘ Best Practice – In cases where the Recursive method causes high load on the directory server and the PDP process, the Nested Groups feature can be disabled or the depth can be configured to the lowest possible value. Use Multi per-group for an Active Directory environment with many defined users and groups with less groups defined in Access Roles.

Policy Enforcement

Once the PDP has the information needed to resolve the Access Role, the information is shared with the Policy Enforcement Points (PEP), which enforce the identity-based policy. This can be on the same or a separate gateway. In large-scale scenarios, you may want to have dedicated PDP instances with separate PEPs.

ⓘ To configure a gateway as a PEP only, enable Identity Awareness, finish the wizard by selecting Cancel. Select Identity Awareness and disable all identity sources. In Identity Awareness -> Identity Sharing enable Get identities from other gateways.

Identity Sharing (sk149255)

Identity Awareness Security Gateways can share the identity information that they acquire with other Identity Awareness Security Gateways. Users who need to pass through many Security Gateways are only identified once, without creating additional load on the identity sources or interfering with a streamlined end-user experience.

ⓘ In a distributed system with multiple Identity Awareness Security Gateways, we recommend using Identity Sharing.

There are multiple ways to deploy Identity Sharing:

- a PDP shares identities to multiple PEPs
- a PEP receives identities from multiple PDPs
- PDP and PEP processes run on different gateways and communicate using a pull sharing method
- PDP and PEP processes run on the same gateway and communicate using a push sharing method.

Push Sharing Method

In this method, the PDP sends or pushes each identity immediately to the PEP as it is acquired. This is the only sharing method used when the PDP and PEP run on the same gateway.

Smart-Pull Sharing Method

In this method, identities are sent to the PEP only when the PEP needs them, i.e. requests or pulls them from the PDP. In larger deployments not all identities acquired by PDPs are needed by all of the PEPs. For instance, small branch offices with a small number of users do not require storing of all of the identities acquired by the PDP located in the headquarters site. Storing unnecessary identities consumes more space on the PEP and creates unnecessary transactions between the PDP and the PEP over the network.

Smart-Pull is a 3-Step Process

1. Identity Acquisition
   Identities are acquired by the PDP and stored in the PDP's repository. The PDP notifies the relevant PEPs that it's aware of the network (Class C) where the user is located, but does not publish the identities to the PEPs.
   To show all of the networks a PDP has published, run the following command on the PDP:
   ```
   # pdp network info
   ```
   On the PEP use this command to show the networks that the PDP is aware of:
   ```
   # pep show network pdp
   ```

2. Sub-network Registration
   When a user initiates a connection through the PEP, where the policy requires an identity element, then the PEP searches for the identity in its local database. If not found, the PEP checks to see if there is a PDP that knows that the Class C network needed to resolve the identity. If found, then the PEP registers to that PDP, asking to be notified about a smaller sub-network, e.g. a mask 255.255.255.240 instead of the larger Class C mask of 255.255.255.255.0.
   On the PEP, to see which 255.255.255.240 network the PEP is registered to, use this command:
   ```
   # pep show network registration
   ```
   On the PDP, to see which PEPs are registered to which 255.255.255.240 sub-networks, use this command:
   ```
   # pdp network registered
   ```
   The PDP will publish all currently known identities from this 255.255.255.240 sub-network to the registering PEP.

3. Identity Propagation
   When the PDP acquires the identity of a user whose IP is in a 255.255.255.240 sub-network that a PEP has registered, the PDP immediately publishes the identity to the registered PEP.

Identity Awareness Role with CloudGuard Controller in Cloud Environments

Identity Awareness also serves another important role with the CloudGuard Controller in our integration with Software-Defined Data Centers (SDDC) and virtual cloud environments. The CloudGuard Controller automatically updates the security policy and security logs as virtual appliances, computers, devices and IP addresses change in these dynamic environments. With the CloudGuard Controller and Identity Awareness, Check Point gateways integrate seamlessly within the following virtual cloud environments: Amazon Web Services (AWS), Microsoft Azure, Cisco ACI, Cisco ISE, Google Cloud Platform (GCP), Nuage Networks VSP, OpenStack, VMware vCenter and VMware NSX.

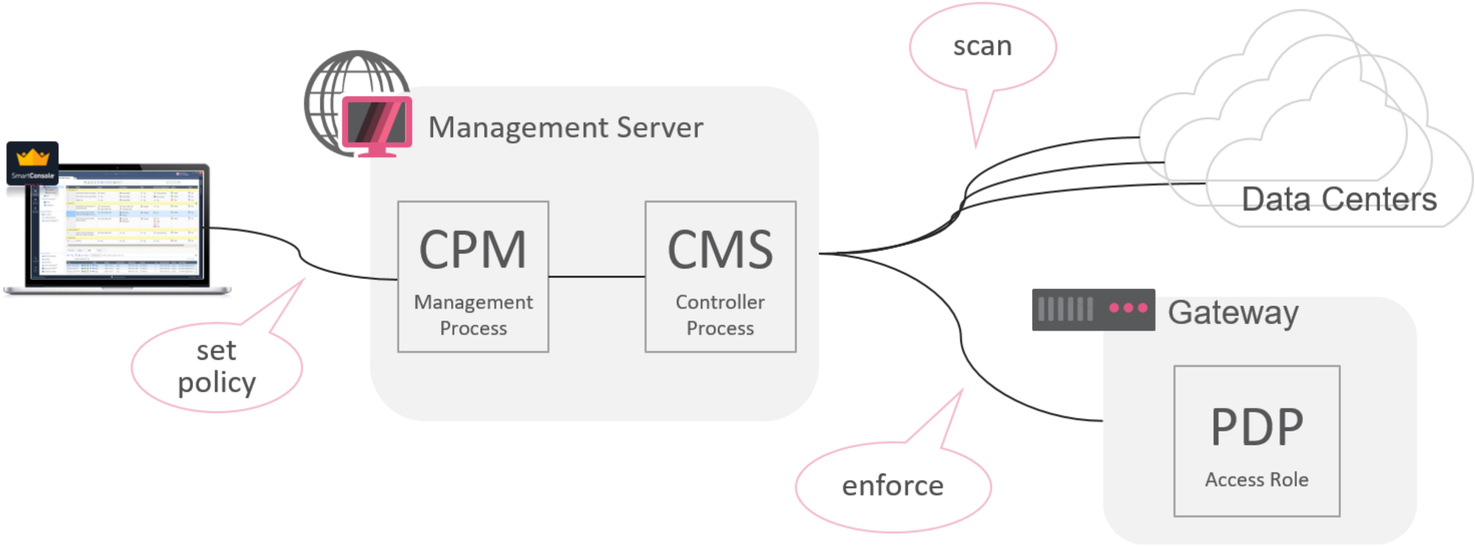

Here's how it works: the CloudGuard Controller scanner periodically polls objects in the data center. The scanner then connects to the management CPM process to update the data center objects and it also connects to the PDP process on the gateways to update the data center objects used in the gateway security policy (sk115657).

Cisco ISE CloudGuard Controller Integration vs Identity Collector Integration

Cisco switches and wireless controllers embedded with Cisco TrustSec technology support the assignment of SGTs. An SGT can be assigned dynamically or statically. Dynamic classification occurs via an authentication sequence, via 802.1x, MAB (MAC Authentication Bypass), or web authentication. When authentication isn't available, static classification methods are necessary. In static classification the tag maps to some identification element (an IP address, a subnet, a VLAN, or an interface) rather than relying on an authorization from Cisco ISE. This process of assigning the SGT is defined as "classification." Static classifications are commonly used for static devices, such as data center servers, or topology based policies, such as a subnet based policy.

Use CloudGuard Controller to enforce SGTs statically mapped to IP addresses. Use Identity Collector "learn" login events generated by Cisco ISE when an SGT is dynamically assigned to users and devices. Do this by creating an Identity Tag object in SmartConsole and then reference this in the Access Role object.

KEY PRINCIPLES OF IDENTITY AWARENESS DESIGN

Every IT organization from large enterprises and small businesses to Services Providers can greatly enhance their security posture and increase the overall value of their network security deployment with contextual identity-based metadata.

- Dynamic user-based access control policies automatically react to changes in users status and logical grouping
- User data enhances security and access logs, allowing for better visibility and enforcement
- Threat prevention and security events analysis can contain user data for easier compliance
- Bandwidth control can be applied to role-based access controls to internal or external assets

The next sections highlight key principles and best practices to follow for designing and implementing a secured and highly integrated user-aware architecture.

Ask the Right Questions

Before entering the design stage, first start with an in-depth discovery process to better understand the customer's topology, users, connected devices, identity sources and identity topology.

Customer environment topology

- How do users access the network?
- At which enforcement point are users inspected?
- Do the branches need to communicate with each other or just with the shared services on the core site?

Users

- How many are corporate or managed users?
- How many are guests?
- Do all of the users work concurrently?

Machines

- BYOD or company device?
- Can you deploy an agent on the machine?
- Do you need the machine context and data in logs and enforcement?

Organization Structure

- How many sub-domains are defined?
- What is the interaction between users in different domains?

Identity Sources

- Which authentication services are available?
- Are there customized authentication solutions?
- Do different types of users or machines use different authentication services?

IDA Topology

- Where will the gateways connected to the ID source be deployed?
- Which gateways will this information be shared with?

AD Related

- What's the number of Active Directory domains?
    - Are they fully trusted among each other?
- How many Logon Servers do we have per location/logical area/business unit?
- What's the average number of groups a user belongs to?
- Are there nested groups?
- Are there cyclic groups where groups are members of each other?

Become Familiar with Your Identity Sources and Integrate Them Correctly

Designating the optimal Identity Sources to integrate with Check Point's security gateways is an essential step of any design process. Modern networks offer a wide variety of possible sources for users and device identities, such as Wi-Fi access points, NAC solutions, Active Directory, a SAML based authentication service and endpoint agents and clients.

Becoming familiar with these sources and actively searching for other potential sources based on their external interfaces can make the difference between an under-optimized solution and a winning design. Consider APIs implemented by Check Point such as Cisco ISE or third party sources such as Pulse Secure, Forescout CounterACT, RADIUS accounting and Syslog.

While it remains common and often recommended to integrate the Check Point deployment with Microsoft Active Directory or other LDAP-based directory services, a well-designed solution takes all of the possible integrated sources into account and uses the sources that are most suitable in terms of simplicity, effectiveness and richness of metadata, performance and resource consumption.

For larger scale networks and/or multi-site deployments, there is no single identity source that can provide this function across all networks. For example, branches of the same enterprise may use different IAM technologies. Connecting the local security gateway to this branch equipment may provide better visibility and performance than connecting to a remote corporate Domain Controller located at headquarters.

Each Identity Source introduces its own pros, cons and system limits, resulting in different capacities for concurrent authenticated sessions, load on the source, load on the Security Gateway, permissions required for integration, and level of contextual data. Use this document, *Identity Awareness Reference Architecture and Best Practices,* the *Identity Awareness Administration Guide*, and relevant SecureKnowledge articles to create a winning design.

Identity Use Cases

As described in the *Identity Awareness Administration Guide*, sometimes the use case dictates which Identity Source to use, e.g. when terminal servers are deployed in the customer environment. These are the priorities of the different Identity Sources. Some are better suited for small (S), medium (M) or large (L) scale deployments.

1. Remote Access
2. Terminal Servers Agent (S), Identity Agent (S)
3. Identity Collector (L), RADIUS Accounting (L), Captive Portal (M), Identity Awareness API (M)
4. AD Query (S)

Also some Identity Sources are complimentary, e.g. Identity Collector or AD Query, which are seamless to the end user are complimentary with the more obtrusive captive portal, which captures devices that don’t login to Active Directory.

⚠ Identity Collector and AD Query should not be used together as they collect from the same identity source.

Requirement	Recommended Identity Source
Users access the organization through VPN	**Remote Access**. Let’s you identify Mobile Access and IPsec VPN clients that work with Office Mode.
Terminal Servers and Citrix environments	**Terminal Servers (S)**. Requires you to install the Terminal Servers Identity Agent on each Terminal Server.
Application Control	**Identity Collector (L) or AD Query (S) and Browser-Based Authentication (M):** The AD Query finds all AD users and computers. The Browser-Based Authentication identity source is necessary to include all non-Windows users. It also serves as a fallback option, if AD Query cannot identify a user. If you configure Transparent Kerberos Authentication, then the browser attempts to authenticate users transparently by getting identity information before the Captive Portal username/password page appears for the user.
Data center or internal server protection	The options are: • **Identity Collector (L) or AD Query (S) and Browser-Based Authentication (M):** When most users are desktop users (not remote users) and easy deployment is important. **Note** - You can add **Identity Agents (S)** if you have mobile users as well as users that are not identified by AD Query. Users that are not identified encounter redirects to the Captive Portal. • **Identity Agents (S) and Browser-Based Authentication (M):** When a high level of security is necessary. The Captive Portal is used for distributing the Identity Agent. IP Spoofing protection can be set to prevent packets from being IP spoofed.
Environment that use a **Cisco ISE** server for authentication, a **NetIQ eDirectory** LDAP server or requires **syslog** parsing	**Identity Collector (L)**: Create Access Roles based upon the Cisco TrustSec network Security Group Tags (SGT).

Identities authenticate to an **external NAC solution**	The options are: **RADIUS Accounting (L)**: Make sure to configure the Security Gateway as a RADIUS Accounting client and give it access permissions and a shared secret. **Identity Awareness API (M)** when supported by the NAC solution. **Identity Collector (L)** using the parsing syslog option.
Logging and auditing with basic enforcement	**AD Query (S)**.

Scale and Impact on the Resource

Nothing stands out more than a system that is not able to handle the number of subscribed users. You may be familiar with security gateways that are undersized. In your Identity architecture design, you'll want to factor in the impact on the Identity Source as well. Below you'll find the recommended deployment size for each connector and the impact of each on the Identity Source. In the discovery phase try to understand what the current load is on these systems so that there are no surprises when your design is implemented. This is sometimes difficult to gauge. In these cases, it can help to recommend a phased roll-out to measure the effect at the end of each stage.

	Terminal Server Agent	Endpoint Identity Agent	AD Query	Captive Portal	IDA Web API	Identity Collector (IDC)	RADIUS Accounting
Deployment Size	Small	Small	Small	Med-high	Med-high	Large	Large
Metrics	Small # of servers/gateway	20,000 users/PDP	800 events/s			1900 events/s, 35 DC per IDC	
Resource Use	High	High	High	Low	Low	Low	Low
Integration Point	AD	AD	AD	AD, RADIUS Authentication	NAC Solutions	AD, Cisco ISE, Syslog, eDirectory	NAC Solutions

Design your Users Network Topology with Identity in Mind

To design a feasible and effective identity and access solution, it's essential to plan and deploy a topology that supports the identity acquisition and enforcement requirements. Start by mapping the user access layers and locations to the network. This simple process can help in spotting missed identity sources and more importantly in closing "identity blind spots", where traffic is generated without any feasible method to correlate it to users or devices.

A common example of a blind spot is a user's access segment NATed by a device that masks the original source IP. Another example is a remote access client, either Check Point or a third party that can provide significant information directly from the endpoint. A third example is a terminal server that connects multiple users. For the initial design stage and especially for existing "brown field" networks, it's important to view the network as a collection of identity sources and authentication points that can be used to increase the effectiveness of identification.

Another factor of the design is the enforcement points. Locating the optimal decision and enforcement points is critical. For example, consider a distributed enterprise network with global branches, all connecting to two data center sites in the cloud. In theory, the decision and enforcement points could be concentrated in the data center perimeter gateways, feeding from a cross-organization directory service. This would not be ideal though, considering the significant logical distance between the users and the enforcement points; NAT, VPN tunnels, multiple hops and interfaces changes, etc.

A much better approach in such cases would be to acquire and enforce the identity and access policy closer to the users on th e branch sites. In this method, each branch gateway could make use of the local identity sources and make decisions based on Identity-enhanced policy to increase effectiveness, save bandwidth and reduce load from the data center sites' gateways and shared services. While many solutions make it cumbersome to achieve, with Check Point's central management architecture the deployment is trivial.

Designating the Policy Decision Points and deploying the Policy Enforcement Points correctly on the optimal gateways, while u sing the Identity Sharing capability when needed is essential. To complete this perspective, consider that the organization infrastructure does not only offer opportunity, it also introduces constraints and limitations.

- *High Latency Connections* such as VPNs between remote sites.
- *Compute Limitations on user's devices* that may be impacted by running another agent.
- *Compute Limitations on a security gateway* that may not be properly sized to process the required number of users.
- *Resource Limitations on the identity sources* such as a heavily utilized Domain Controller that cannot allocate more resources to a Check Point connector.
- *Limited Permissions and Collaboration* between IT departments that prevents mutual interfaces and assigning permissions to third party platforms, or distributing agents via group policies, etc.

Some constraints may also originate from the Check Point deployment itself. For example, when a Multi-Domain management server with multiple Domain Managers needs to be integrated with the same sources or share identities between their managed Security Gateways. Those constraints must also be addressed.

Understand Identity and Access Flows in Your Organization

There is simply no single solution or architecture that fits every organization and use case. Understanding users' access to the network and the resources they need to access is the third key for a successful design.

Once the identity sources have been designated and the enforcement points decided, the next step is to understand the authentication flow of users with the identity source, the structure and metadata in which the users are mapped and the desired connections for authenticated users. Even more important are the expected user experience and constraints – not every deployment is compatible with endpoint clients and not every group of users can be asked to use authentication portals. The methods and flow of connecting and authenticating users is the final and most critical concept.

Security requirements often obstruct a streamlined user experience, but a well-designed solution should combine those factors organically – whether by integrating with Active Directory to retrieve login events and using them to map IPs to users that need to authenticate and login regardless of the network security solution, or by using a RADIUS integration or API integration with NAC solutions that are deployed as a mandatory gate to gain access to the network or by applying any other method supported by Identity Awareness. Providing a stable, frictionless experience to the end-users, with minimal impact on the infrastructure, without compromising on security granularity or creating "identity blind spots" is the ultimate goal for any design.

	Integration Point	End-user Experience	Admin Experience	Identity Assurance	Records Logoff Events
RADIUS Accounting	NAC Solutions	Transparent	Clientless, Easy to apply	Real Time	Yes
Identity Collector	Active Directory, Cisco ISE, Syslog, eDirectory	Transparent	Configure agent on Windows server	Real Time	Yes, for Cisco ISE and syslog
IDA Web API	NAC Solutions	Transparent	A one-time API implementation	Real Time	Yes

Captive Portal	AD, RADIUS Authentication	Might require manual authentication, BYOD, manual authentication in non-Kerberos (SSO)	Easy to apply		Can be configured to record logoff event
Active Directory Query	Active Directory	Transparent	Clientless Easy to apply		
Terminal Server Agent	Active Directory	Transparent	Requires implementation	Very Secure	Yes
Identity Agent	Active Directory	May require user implementation and user interaction	May require admin implementation	Very Secure	Yes

EXAMPLE: DISTRIBUTED ENTERPRISE WITH BRANCH OFFICES

Now that we have gone over some of the principles of designing a solution, consider this example customer environment. What identity sources would you use at headquarters or at the branch sites? Which Check Point connector would provide the best user experience and handle the number of subscribed users? Where would you locate the PDP and PEP in this design? Would you use identity sharing?

Customer Environment	Users	Machines	Organizational Structure	Identity Sources	IDA Topology
Distributed organization with multiple branch sites and an HQ site 5 AD DCs in HQ and connected to IDC	10,000 total 5,000 in HQ of which most are connected via desktops and 50 via Terminal Servers 5,000 across multiple branches	Desktops and Terminal Servers in HQ BYOD mobile devices in branches	Branches do not require connectivity to other branches, only to HQ HQ users do not access branch resources except for a limited group of administrators	Branches use a ForeScout NAC solution that supports RADIUS accounting AD is deployed and contains the entire organization	HQ perimeter gateway is under heavy load LAN gateways are underutilized

Recommended Solution

Headquarters

- Terminal Server Agent
- Identity Collector
- PDP enabled on Internal firewall, shares to perimeter
- PEP enabled on perimeter

Branch

- RADIUS Accounting from ForeScout

COMMON DESIGN MISTAKES AND BEST PRACTICES

Check Point has thousands of customers worldwide and Identity Awareness has been available in Check Point's integrated products since 2010. Let's consider some common design mistakes that we have seen and best practices that we have learned over the years.

Sticking to one Identity Source – for instance integrating every security gateway with all AD Domain Controllers is NOT a recommended practice.

Defining Access Role for Each User – Access Roles are designed to be used as logical containers for groups of users and/or devices that meets certain criteria. As mentioned earlier, using groups in your security policy simplifies your policy, while also creating a more dynamic policy. Adding the user to a group in Active Directory is all that's needed to enforce your group-based policy.

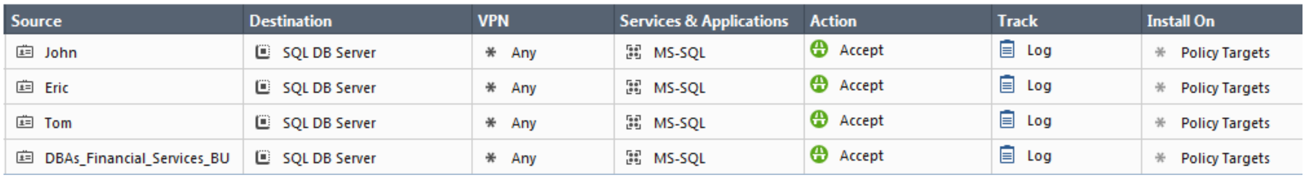

Source	Destination	VPN	Services & Applications	Action	Track	Install On
John	SQL DB Server	Any	MS-SQL	Accept	Log	Policy Targets
Eric	SQL DB Server	Any	MS-SQL	Accept	Log	Policy Targets
Tom	SQL DB Server	Any	MS-SQL	Accept	Log	Policy Targets
DBAs_Financial_Services_BU	SQL DB Server	Any	MS-SQL	Accept	Log	Policy Targets

IAM Mechanisms such as Identity Awareness operate best in a Whitelist Approach, meaning they are designed to grant access based on identity to authenticated users and devices.

Activating Conflicting Identity Sources on the Same Security Gateway – AD Query and Identity Collector conflict and should not be used as the identity connector for the same gateway. Events may arrive out of sync and the same event may be observed multiple times, leading to unpredictable results.

Activating All Gateways to Get Identities from the Same Source - If more than one Security Gateway gets identities from the same AD server, the AD server can become overloaded with WMI queries. In environments that use many Security Gateways and AD Query or Identity Collector, we recommend setting only one Security Gateway to acquire identities from a given Active Directory domain controller for each physical site and sharing identities with the other gateways.

Meshing Identity Sharing - It's not advisable to activate identity sharing on every security gateway in the network and mesh it with all of the other security gateways. For instance, in a fully Trusted AD network you shouldn´t do PDP to PEP sharing. The same information will be shared within the trusted AD network and then arrive "again" when shared from the PDPs to the PEP instances.

Ignoring Network and Services Constraints – This can result in a serious impact on services and infrastructures. For example, integrating an already loaded Domain Controller via AD Query to a large number of Security Gateways could cause it to become unresponsive.

Forgetting to Exclude Services (see sk131792) – When using AD Query it is highly recommended to activate "assume only one user per device" or Identity Collector which "assumes one user per device" by default and exclude any non-user devices that may be inspected, such as Exchange servers or Citrix servers. It's also highly recommended to exclude all known service accounts. These are not used in the user-based policy and so they create an unnecessary overhead.

Forgetting to Exclude Multi-user Hosts – When using ADQuery or Identity Collector

Ignoring Check Point Limitations – Do not ignore PDP or PEP process limits. For example, the PDP and PEP limit on R80.10 and up is up to 200,000 associated IPs. Identity agents are limited to 20,000 per PDP. This limit is not related to PEP.

Nested Groups (see sk128212) – Be aware of limitations and resource constraints of the three nested group's options. See the best practice section above. Enabling nested groups may cause high load on CPU and unresponsiveness of the PDP daemon, as well as load on the Domain Controllers, particularly during policy installation.

Identity Agents Limitations – Identity Agents work well in small deployments, i.e. less than 20,000 users per PDP. By selecting which gateway the Identity Agent connects to, you can manage the load. For a branch office or an organization with one or more gateways protecting a single data center and another perimeter Security Gateway, configure the agents to connect to the one Security Gateway and then share the identity with the other Security Gateways in the network. For complex multi-data center environments with several Security Gateways that protect different data centers and the perimeter, we recommend balancing the Identity Agents' authentication with different Security Gateways. See the *Identity Awareness Admin Guide*.

Identity Agent Keep-Alives and Connection Retries – An unstable network or high utilization of the Identity Awareness daemon, such as during policy installation, may interfere with the Identity Awareness daemon's responsiveness for a short while, causing Identity Agents to become disconnected if the user sessions disconnect.

To avoid this, the Security Gateway and the Identity Agent client can be configured to be more tolerant of various connection issues. In large environments, the default value of 5 minutes between keep-alive messages can cause the number of open file descriptors on the Security Gateway to reach the default limit of 1024. One solution is to increase the interval between keep-alive messages to 10 minutes. Another solution is to modify the number of attempts the Identity Agent tries to connect, so that a temporary connectivity shortage will not cause a client disconnection. See the *Identity Awareness Admin Guide*.

Overlapping Identities – Customer identity sources may contain the same users on different Identity Sources, causing unpredictable results. When searching for a user object, the gateway first checks the internal Check Point user database. If no match is found, it sends a query to all LDAP Account Units simultaneously looking for the user object. The first AU that responds is the one the user is matched to. If your LDAP groups used in the rules are tied to the other LDAP AU, then the user will not match. The connection will fall down to the cleanup rule and be dropped.

SUMMARY

Every IT organization can greatly enhance their security posture and increase the overall value of their network security deployment with contextual identity-based metadata. With our Identity Awareness API and integrations with leading directory stores and IAM vendors, Check Point security gateways fit well in large and small environments.

To create a winning design that is satisfying to your customer's end users and IT staff, start with an in-depth discovery process to better understand the customer's topology, users, connected devices, identity sources and identity topology. Ask the right questions. Become familiar with all possible identity sources. Design your Check Point deployment with identity in mind. Understand identity, access flows and use cases. Understand the impact on users, on the identity sources and the Check Point infrastructure.

As a trusted Check Point expert, understand our product. Read and understand the Check Point User and Administrator guides. Complete our eLearning modules. Find questions to common problems in Check Point SecureKnowledge articles and Advanced Technical Reference Guides. Consider and learn from some of the common design mistakes that we have seen and best practices we have learned over the years that are outlined in this guide.

If you get stuck, help is available. Customers can get help directly from our Check Point *Technical Assistance Center* and *Partners*. For more complicated and larger scale deployments, *Professional Services* is ready to help.

Finally start a discussion, ask a question and share your experience on Check Point *CheckMates User Community*.

REFERENCES

[1] Forrester Research, Evolve Your IAM Strategy for Your Digital Business, July 20 2018, https://www.forrester.com/report/Evolve+Your+IAM+Strategy+For+Your+Digital+Business/-/E-RES81861

APPENDIX A: SECUREKNOWLEDGE ARTICLES

Title		Link
AD Query	ATRG: Understand how AD Query works	sk86441
AD Logon Servers	Best Practices Large Scale Deployment	sk88520
Identity Collector	Technical Overview	sk108325
CloudGuard Controller	ATRG: CloudGuard Controller	sk115657
Trust, Cisco ISE to IDC	Trust is based on certificates, follow sk114436 to use self-signed certificates	sk114436
SmartConsole and AD	Understand the communication between SmartConsole and AD Servers	sk115677
SmartConsole and AD	Secure the communication between SmartConsole and AD Server using Kerberos	sk111996
Kerberos	Troubleshooting Kerberos in Identity Awareness	sk104055
Kerberos Ticket Size	In many environments the Kerberos tickets are larger than the default value defined. Increase the ccc_max_msg_size to the maximum value of 65523	sk66087
SSO Security	AES encryption types missing from SSO configuration	sk111945
ID Agent for Mac	Identity Awareness Agent for MacOS	sk63920
ID Agent Details	Identity Awareness Agent Network Communication and Processes	sk11323
Nested Groups	Nested Groups Improved Capabilities	sk128212
Nested Groups	Controlling Nested Groups	k66561
AD Service Accounts	Exclude Service Accounts	sk131792
Identity Sharing		sk149255

APPENDIX B: HOW IT WORKS

Remote Access

Identities are acquired for **Mobile Access** clients and **IPsec VPN** clients configured to work in Office Mode when they connect to the Security Gateway.

Session Details

- IP, User, remote access client

Authentication Process

- On access acquired from Internal, AD, RADIUS, TACACS, SecurID, Certificate

Browser-Based Authentication

Identities are acquired through authentication **web captive portal** on Identity Awareness Gateway.

Use Cases

- Guest authentication
- Backup if Transparent Kerberos Authentication fails
- Option for user to install of light Identity Agent

Session Details

- IP, User

Authentication Process

- On access acquired from Internal database, AD, RADIUS

Browser-Based Authentication Single Sign-On

The transparent Kerberos Authentication Single-Sign On (SSO) solution transparently authenticates users already logged into the AD. When a user authenticates to the domain, the user gets access to all authorized network resources and does not have to enter credentials again. If Transparent Kerberos Authentication fails, the user is redirected to the Captive Portal for manual auth entication.

Identity Agents

Identities are acquired using full, light or custom configured **endpoint agents** that are installed on the Endpoint computers.

Use Case

- High level of security; packet tagging to prevent IP spoofing, IP change detection
- Transparent authentication with Kerberos Single Sign-On
- Connectivity even while roaming to another network

Session Details

- IP, User, AD Machine in an Active Directory environment

Authentication Process

- On access acquired from Internal, AD, Kerberos Single Sign-On

Identity Agent Kerberos Single Sign-On

(Identity Awareness Administration Guide) The Kerberos protocol is based on the idea of *tickets*, encrypted data packets issued by a trusted authority, which in this case, is the Active Directory (AD). When a user logs in, the user authenticates to a domain controller that provides an initial *ticket granting ticket* (TGT). This ticket vouches for the user's identity. When the user needs to authenticate against the Identity Awareness Gateway, the Identity Agent presents this ticket to the domain controller and requests a *service ticket* (SR) for a specific resource (Security Gateway that Identity Agents connect to). The Identity Agent then presents this service ticket to the Security Gateway that grants access.

How SSO Works

This is the workflow for SSO (Single Sign On):

1. The user logs in to the computer and authenticates to the AD server.
2. The AD sends an initial ticket (TGT) to the computer.
3. The Identity Agent connects to the Security Gateway, which then requests the identity.
4. The Identity Agent requests an SR (service ticket) for the Security Gateway and presents the TGT to the AD server.
5. The AD server sends the SR to the computer.

The user name is encrypted with the shared secret between the Security Gateway and the AD server.

6. The Identity Agent sends the SR to the Security Gateway.
7. The Security Gateway uses the shared secret to decrypt the ticket and confirms the user identity.
8. The user can access the Data Center.

Custom Identity Agents

You can use the Identity Awareness Configuration Utility to create custom Identity Agent installation packages (the Identity Awareness Configuration Utility - IAConfigTool.exe - is installed as part of Identity Agent). Identity Agents have many advanced configuration parameters. Some of these parameters are related to the installation process, while others are related to Identity Agent functionality. All of the configuration parameters have default values that are deployed with the product and can remain unchanged.

Identity Agent Type	Description
Full	Predefined Identity Agent that includes packet tagging and computer authentication. It applies to all users of the computer on which it is installed. Administrator permissions are required to use the Full Identity Agent type.
Light	Predefined Identity Agent that does not include packet tagging and computer authentication. You can install this Identity Agent individually for each user on the target computer. Administrator permissions are not required.
Terminal Servers	Predefined Identity Agent that installs Managed Asset Detection (MAD) services and the Multi-user host driver on Citrix and Terminal Servers. This Identity Agent type cannot be used for endpoint computers.
Custom	Let's you configure custom features for all computers that use this agent, such as MAD services and packet tagging.

Custom Features

Select these features for the Custom Identity Agent type:

- **MAD Service** - Install MAD (Managed Asset Detection) services for Kerberos SSO and computer authentication.
- **Packet Tagging** - Install the packet tagging driver to enable anti-spoofing protection. The driver signs every packet that is sent from the computer. This setting is required if you have Firewall rules that use **Access Roles.** IP Spoofing is enabled.

Terminal Servers

Identities are acquired using agents that are installed on a **Windows-based application server that hosts Terminal Servers, Citrix XenApp, and Citrix Xen Desktop services**. These agents are used to identify individual user traffic coming from Terminal Servers.

Use Case

- Any network where terminal servers are used

Session Details

- IP, AD User

To explain it simply, the Terminal Server Identity Agent that is installed on the Terminal Server communicates to the Security Gateway about how it will control the connections for each user (explained below). This information is later used when the traffic reaches the Identity Gateway.

The Terminal Server Agent communicates with the gateway over SSL (usually port 443 unless configured differently).

The solution is based on source ports. The Terminal Server Identity Agent installs a TDI driver that intercepts all requests from any process that requests a new connection. Once the request reaches the TDI driver, it queries the system to fetch the requesting user behind this new connection and chooses a source port from a pool of port ranges allocated for this specific user.

Two different users will have two different port range pools, thus allowing the Identity Gateway to distinguish between the different connection owners.

The solution does not support non-port based protocols. The solution supports TCP and UDP protocols only.

Active Directoy Integration Comparison

Logon Event Learning Method	Required User Privileges	Comment
AD Query (WMI)	Server Operator	High privileged account
Identity Collector	Event Log Reader	Low privileged account
LDAP Query as configured in the LDAP Account Unit	Domain User	Very low privileged account

Active Directory Query (sk60301)

Identities are acquired seamlessly from **Microsoft Active Directory**.

Use Cases

- Active Directory environments

Session Details

- IP, AD User, AD Machine, AD group

The steps in the AD Query example are listed below:

- The Security Gateway registers to receive Security Event logs from the Active Directory Domain Controllers.
- A user logs in to a desktop computer using his Active Directory credentials.

- The user logs on to a Domain Controller (DC). This DC and only it will have the logon security log (those logs do not replicate to other DCs. The firewall queries all of the DCs security logs and gets the logon security log.
- The Active Directory DC sends the Security Event log to the Security Gateway. The Security Gateway extracts the user and IP address information (user name@domain, machine name and source IP address).
- The user initiates a connection to the Internet.
- The Security Gateway confirms that the user has been identified and allows him to access the Internet, based on the policy.

AD Query (ADQ) employs the following four main steps:

- Communication with the Domain Controllers and subscription for the relevant Security Event logs.
- Extracting the user and/or machine to IP Association from the Event Log.
- Filtering undesirable associations.
- Applying the new Association with the user/machine to the IP database.

ⓘ **Best Practice**: In environments that use many Security Gateways and AD Query, we recommend that you set only one Security Gateway to acquire identities from a given Active Directory domain controller for each physical site. If more than one Securi ty Gateway gets identities from the same AD server, the AD server can become overloaded with WMI queries.
Set these options on the **Identity Awareness** > **Identity Sharing** page of the Security Gateway object:

- One Security Gateway to *share* identities with other Security Gateways. This is the Security Gateway that gets identities from a given Domain Controller.
- All other Security Gateways to *get* identities from the Security Gateway that acquires identities from the given Domain Controller.

Identity Collector (sk108235)

Identities are acquired using agents that are installed on a central Windows server. The agents communicate with both **Microsoft Active Directory** Domain Controllers, **Cisco ISE** servers, **NetIQ eDirectory** LDAP server and are able to parse **syslog messages** to extract identities from it.

Use Cases

- Any network where one of the above is deployed

Session Details

- IP, User, Machine, Cisco Security Group Tag (SGT)

RADIUS Accounting

Identities are acquired using RADIUS accounting directly from a RADIUS accounting client.

Use Cases

- Any network where a RADIUS server is deployed

Session Details

- IP, User, Machine, Group

Identity Web API

Gives you a flexible method for creating identities and easily perform third party integrations

Use Cases

- Any network where a supported third party integration is deployed
- Automation of identity tasks
- Network Access Control (NAC).e.g. quarantine of infected hosts

Session Details

- IP, User, Machine, machine group, fetch machine group from directory, domain, user group, fetch user group directory, role, IA defines role, machine OS, host type