User Identification Operation

PANOS 4.0

Nick Piagentini

Contents

PANOS 4.0 User ID functions 1
User / Group Enumeration 1
1. Using AD User Agent for Enumeration 2
2. Using LDAP Servers for Enumeration 3
User to IP Mapping 5
1. AD UserAgent 5
2. LDAP User Agent 7
3. Captive Portal 8
4. Terminal Server Agent 12
5. Palo Alto Networks client software 12

PANOS 4.0 User ID functions

User Identification in PANOS 4.0 encompasses two primary functions:

- Enumeration of users and their associated group membership
- Mapping of those users to their current IP addresses .

Each of these functions can be performed by different methods. Some methods are effective in specific network environments and some are applicable in all environments.. Both components will be discussed in this document.

User / Group Enumeration

Before a security policy can be written for groups of users, the relationships between the users and the groups they are members of must be established. This information is retrieved from an LDAP directory, such as Active Directory or eDirectory. The firewall or an agent will access the directory and search for group objects. Each group object will contain a list of user objects that are members. This list will be evaluated and will become the list of users and groups available in security policy and authentication profiles. There are two methods for retrieving this data:

1. use an agent that talks to Active directory, or
2. use an agent that talks to LDAP servers.

Both of these methods are discussed below.

1. Using AD User Agent for Enumeration

Operation:

This agent is installed as a windows service on a Windows server that is a member of the domain to be polled. It is configured with a list of Domain Controllers in a single Windows Domain, and will access the first DC on that list for user and group information. If the first DC is not available when it does group enumeration, the agent will continue down the list until it locates a DC that is available.

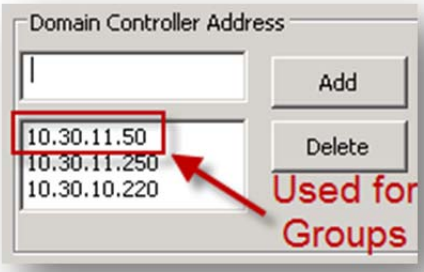

The agent will access the domain controller using Microsoft RPC and will read all of the security groups in the domain. Since the agent is only configured to map users from a single domain, any accounts from other domains will be ignored. For this reason it is a best practice to build security policy using Domain Global groups, as they will only contain users from a single domain and will be correctly represented by the AD user Agent.

After the agent has parsed the domain for groups and users, it can apply a group filter to send only select groups to the firewall. It is strongly recommended that you configure a group filter. By eliminating unneeded groups from the list that is sent to the firewall, overall processing on the firewall's Management Plane is reduced, and the group selection interface in the UI is more succinct/user friendly.

After the initial group membership is obtained, the agent will check to see if group membership has changed every so often, based upon a configurable timer (called "User Membership Timer"). The agent will update the firewall with only the groups that have changed membership. If no changes to group membership is detected there will be no data sent to the firewall.

A single agent can only monitor domain controllers from a single domain. The agent can monitor up to 100 individual DC's from that single domain. In a multi domain environment there will need to be multiple agents deployed, so that group information can be retrieved from all the domains.

For each domain, the firewall (or Virtual System if the firewall is operating in that mode) will select a single agent to gather group data from. By default it will be the first agent configured, but if that agent is not available the firewall will try other agents in the list. To determine the agent being used by the firewall for group membership the >show user pan-agent statistics command can be used. The agent with the * before the word *connected* is the one being used for all group membership, as seen in the screen shot below:

```
---------------- --------------- ----- ------- ------------------ ------ ------
-------- -------- -------- --------------- ---
Pan Agent        192.168.100.12  3033  vsys1   *connected, ok     4403   49
4095     1462733  600      pancademo       1
```

Best Practices for AD User Agent

1.) Configure well connected Domain Controllers on the top of the list in the agent and firewall configuration.
2.) Filter the list of groups that is sent to the firewall to include only the groups that will be used in firewall policy. If you want to make sure that all users are tracked, include the group "Domain Users".
3.) Only use Domain Global groups in firewall policy when operating in a multi domain environment. Note that this is not in line with traditional Microsoft AD practice, where Domain Local groups are used to control rights and access.
4.) If some agents are located across slow or heavily impacted links it may be best to configure only the well-connected agents first and run a commit. This will get the initial users and groups on the firewall and insure that future updates are just deltas.

2. Using LDAP Servers for Enumeration

Operation:

The Palo Alto Networks next generation firewall can gather user and group information from an LDAP directory without the use of an agent. This method can be used to enumerate Active Directory or any other LDAP environment. The firewall will perform all of the LDAP connections and no agent is required for this function.

The firewall defines a number of "LDAP Servers" under the User Identification node. Each LDAP Server instance represents a bind to a specific part of an LDAP tree. It will enumerate all of the user and group objects at that point and below. Filters can be defined in this configuration using standard LDAP syntax to limit the users and groups returned. If this method is to be used to enumerate users from Active Directory, there will need to be a LDAP Server configured for each domain. Global Catalogs cannot be used for user and group enumeration across AD domains. Only LDAP objects that use a field to list membership can be used as groups on the firewall. PANOS does not support the use of container objects such as Organizational Units (OU) as security principals in firewall policy.

Access credentials to the LDAP tree is specified in a LDAP Authentication server object that is referenced by

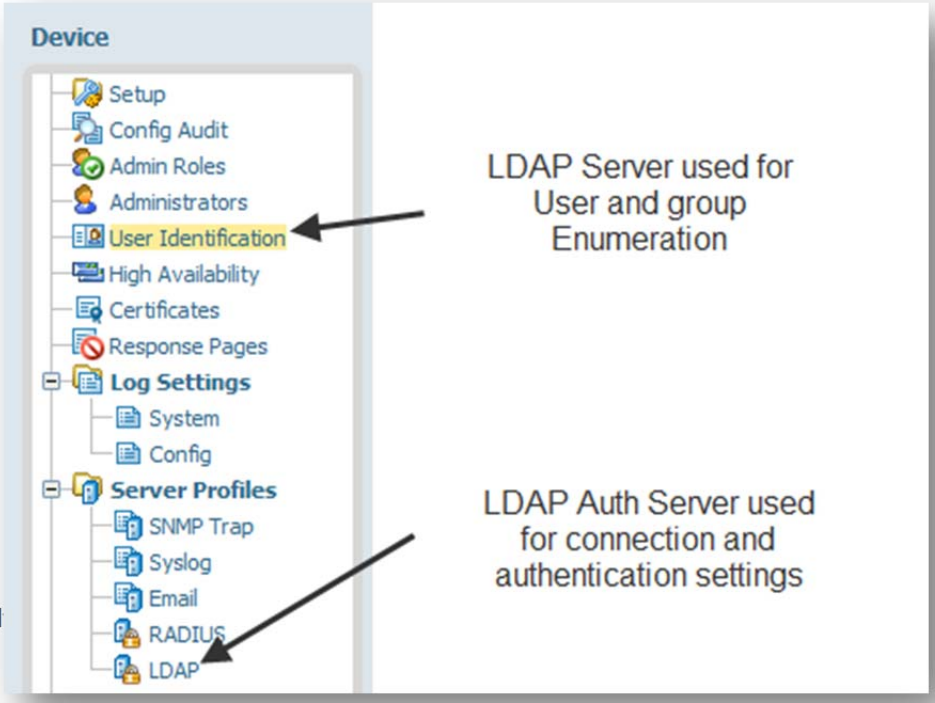

the LDAP Server object. The Authentication Server object also specifies which directory servers will be contacted, the order they will be contacted in and when the firewall will try the next one on the list.

Configuration of the LDAP Server object requires knowledge of the LDAP structure in use, such as types of objects used as groups and users. For example in a standard Active Directory deployment the users are objects "objectclass = User" and are most commonly referred to by either the SAMAccountName

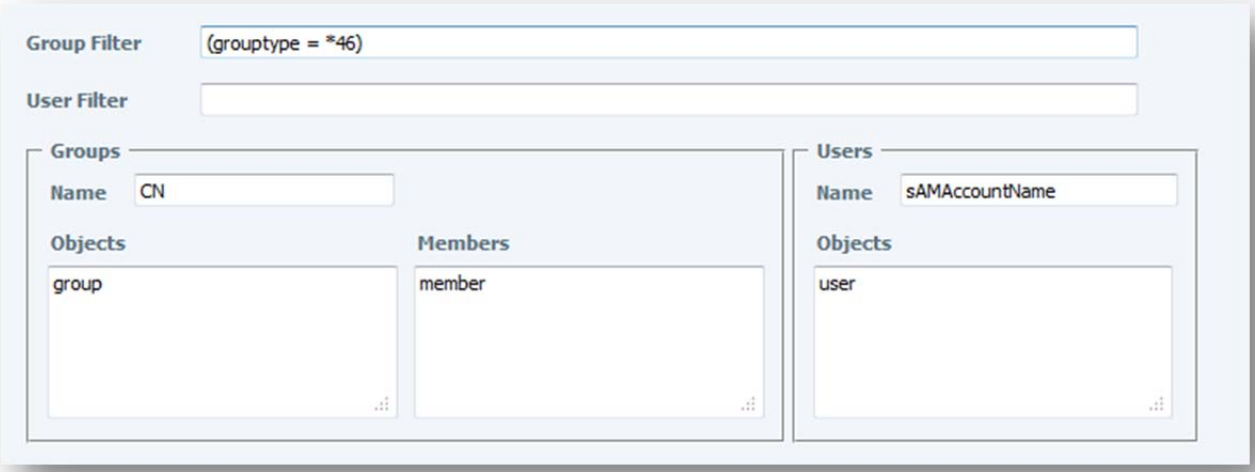

(jdoe) or UserPrincipalName (jdoe@corp.local) fields. The groups " objectclass = group" are referred to by the CN field and store a list of users in a "members" field. This level of information is required to configure the LDAP Server. The following is an example of LDAP server configuration to enumerate users from all Domain Global security groups on an Active Directory domain.

For interoperability between the LDAP server and the AD Agent, a domain can be specified in the server configuration. This domain will be added as a prefix to any user accounts learned by the agent. By synchronizing this value with the NETBIOS name of the AD domain in use, we can map users authenticated by NTLM to users enumerated by LDAP.

Best Practices using LDAP Servers

1.) If the underlying directory is Active Directory, make sure the "Domain" field of the LDAP Server matches the NETBIOS name of the domain.
2.) Use of an LDAP browser can be extremely helpful if working with a non-generic LDAP deployment.
3.) Use group filters to minimize the number of groups returned. For example (grouptype = *46) will return only Domain Global security groups.

User to IP Mapping

The process of mapping users to IP addresses is the more complex of the two User ID tasks. PANOS 4.0 provides multiple methods to map users to IP addresses. Some methods require specific directory structures to be in place. Some methods require software agents or clients to be installed. If any of the methods map a user to an IP address, that data can be used by the firewall in both policy and reporting. User data is written to all appropriate logs when the logs are generated. The methods mapping users to IP are:

1.) AD User Agent
2.) LDAPUser Agent
3.) Captive Portal
4.) Terminal Services Agent
5.) Palo Alto Networks client software (SSL VPN, Global Protect)

Each of these is described below.

1. AD UserAgent

The AD User Agent performs both the enumeration and mapping tasks. Even thought the two processes are separate, the agent cannot be configured to perform only one or the other. In Active Directory environments, the AD Agent is very useful for mapping users and as a result is also commonly used to enumerate users as well. The agent can map users by monitoring events in the security log and by querying end points. These mappings can be re-confirmed by monitoring user connections to the domain controller during the course of work. The firewall

Security Log Reading

The AD Agent will connect to each domain controller in its list and monitor the security log. On the initial connection the agent will read the last 50,000 log entries. After the initial connection, the agent will then monitor all new events. The AD Agent looks for any of the following Microsoft event IDs:

On Windows 2003 DCs :

- 672(Authentication Ticket Granted, which occurs on the logon moment),
- 673(Service Ticket Granted)
- 674(Ticket Granted Renewed which may happen several times during the logon session)

On Windows 2008 DCs:

- 4768(Authentication Ticket Granted)
- 4769(Service Ticket Granted)
- 4770(Ticket Granted Renewed)

These events will contain a user and IP address. The user's domain will be compared to the domain that the agent has been configured to monitor. Users from domains other than the monitored domain will be ignored. Also monitored will be the IP ranges of the users; only Allowed IP ranges (as configured on the

AD agent) will be recorded. Once the username to IP mapping table is created, the agent will send this data to the firewall. The default timing for checking new log events is every second, but this timer is configurable. Note that these events will only be present in the security log if the AD domain is configured to log successful Account Logon events.

Security log reading is low overhead for the Domain Controller and a highly effective method of mapping users in a Microsoft environment. The mappings will be maintained for a configurable time out, which is recommended to be set to half the DHCP lease time used in the environment. Client systems in an AD domain using the default configuration will attempt to renew their tickets every 10 hours.

WMI / NetBIOS Probes

Where the log reading is effectively a passive method of user mapping, probing is an active method. On a configurable interval, the AD Agent will send a probe to each learned IP address in its list to verify that the same user is still logged in. The results of the probe can be used to update the record on the agent and then be passed on to the firewall. Each learned IP will be probed one per interval period. Care should be taken to make sure that large environments have a long enough interval for all IP's to be probed. For example in a network with 6,000 users and an interval of 10 minutes, that would require 10 WMI requests a second from each agent. These probes are queued and processed by the agent as needed.

In addition, when the firewall receives traffic on an interface in a zone with User Identification enabled that is from an IP address that has no user data associated with it, the firewall will send the IP to all the AD gents configured and ask them to probe it to determine the user. This request will be added to the queue along with the known IP addresses waiting to be polled. If the Agent is able to determine the user at the IP based on the probe the information will be sent back to the firewall.

If the WMI or NetBIOS probe fails the IP address will not be probed again until the firewall sees more traffic from it.

NetBIOS probes have no authentication and do not require any specific group membership of the Agent account. A drawback to NetBIOS is that it is not very reliable across larger networks; it is commonly blocked by host based firewalls and will not work for certain modern operating systems. (Anything with NetBIOS over TCP disabled)

WMI queries are far more reliable and are secured by either NTLM or Kerberos based authentication. To perform these queries successfully the agent account needs the rights to read the CIMV2 namespace on the client systems. By default on Domain Administrators have this right. The underlying WMI query that is sent can be simulated with the following command, where *remotecomputer* would be the IP address of the system being probed:

```
wmic /node:remotecomputer computersystem get username
```

Open Server Sessions

Any connections to a file or print service on the Domain Controller will also be read by the agent. If the user / IP combination for the session does not match the combination that the Agent last learned the mapping will be removed and the user at the IP address will become unknown. The agent will not update user data as a result of information learned from the open server sessions. If the open session confirms the user at the IP address then that mapping will have its lifetime renewed.

In the normal operations of an AD domain, users on Windows systems will connect to the *sysvol* share on the domain controller to check for new Group Policy Objects. The default timing for this is 90 minutes with a +/- 30 minute offset. For users connected to the network during a regular work day this process will insure that they remain mapped throughout the day.

Agent and Firewall Communication

Settings on the Agent control how often the agent communicates with the Domain Controllers and hosts on the network (for polling). The firewall has specific, non-configurable timers for its communication to the agent.

- 2 seconds: Get list of new IP / user mapping from agent. This is a delta of new mapping only.
- 2 seconds: Send list of unknown IP addresses that were encountered in traffic to the agent.
- 5 seconds: Get agent status. This is a heartbeat uses to determine the status of each configured agent.
- 10 minutes: Get group membership changes from agent. This is just the delta of changes since the last check.
- 1 hour: Get full list of IP / user mappings from agent

Best Practices for AD Agent:

1) Set the age out timer for the agent to a value close to half the DHCP lease time.
2) Use WMI over NetBIOS if possible.
3) Make sure to plan the interval for probing based on the total number of users in the environment.

2. LDAP User Agent

The LDAP agent provides two very specific functions. One is to access an eDirectory tree and read the logged in IP for each user. When the user logs in to eDirectory, the IP address of the end point is stored in the directory as a field in the user object. This serves a similar function as the AD Agent's log scraping and only works with eDirectory.

The second function of the LDAP agent is to receive XML user information from external sources. This information can both add and remove user IP mappings. Some examples of the API are:

1) Visual Basic based log in and log out scripts that add and remove the user and all the IP addresses of the end station.
2) Perl based scripts for Mac based systems to register users on log in.
3) Modules for NAC appliances that pass on user and IP information to the firewall.

The API passes the data over SSL using a simple XLM format as follows:

```
<uid-message>
     <version>1.0</version>
<type>update</type>
     <payload>
          <login>
               <entry name="domain\uid1" ip="10.1.1.1"/>
               <entry name="domain\uid2" ip="10.1.1.2"/>
               <entry name="domain\uid3" ip="10.1.1.3"/>
          </login>
          <logout>
               <entry name="domain\uid4" ip="10.1.1.4"/>
          </logout>
     </payload>
</uid-message>
```

3. Captive Portal

Captive Portal is an identification method that is not invoked unless there is no user information for HTTP based traffic that the firewall encounters. If a user has been mapped by one of the other possible methods, captive portal will *not* be triggered. Captive portal is traditionally used to identify users that have slipped through the other methods or for environments where the other methods are not appropriate. Captive portal will only be triggered by a session that matches the following criteria:

1) There is no user data for the source IP of the session
2) The session is HTTP traffic
3) The session matches a Captive Portal policy on the firewall

When captive portal is triggered the browser session is interrupted by the firewall and user credentials are requested. Once the user is identified they will remain mapped until either an Idle or hard time out is reached. At that point the user mapping is removed and captive portal may be triggered again.

For firewalls deployed in L2 or Virtual Wire mode captive portal must be configured transparently. In this configuration the firewall will spoof the destination address for use in authentication. This can generate certificate errors if the original communication was over SSL. A more flexible method is a redirect captive portal, where the firewall uses a 302 HTTP error code to redirect the user to a L3 interface owned by the firewall. When using redirect captive portal a specific SSL certificate can be installed for the portal to mitigate any certificate warnings. In addition redirect captive portal can use cookies to mark the session. This will allow the session to remain mapped even after the time outs have expired. Finally redirect captive portal with cookies can support a user that roams from one IP address to another while keeping the session open. When possible, captive portal should always be deployed in redirect mode.

There are three methods for the firewall to extract user data from the browser:

1.) NTLM Authentication
2.) Web Form Captive Portal
3.) Certificate-based Authentication

NTLM Authentication

Microsoft clients can participate in a NTLM challenge and response exchange that consists of 3 messages. The browser will use the credentials of the currently signed in user. Internet Explorer will do this be default, and Firefox can be configured to do this for specific URI's. (In the about:config set the network.automatic-ntlm-auth.trusted-uris vale to the captive portal URI) This authentication is transparent to the user. The user name captured from this method is the NetBIOS name in the form of DOMAIN\USER, it will be mapped to the appropriate user ID if the AD Agent is in use, or if the LDAP Server configured to read the AD domain has the correct value in the *domain* field. If the browser or operating system does not support NTLM authentication, the firewall will fall back to the next form of Captive Portal. When configuring NTLM based authentication for Captive Portal a host name must be provided. For NTLM to work, this host name must *not* be fully qualified. For example, if the DNS name of the portal is **portal1.company.com**, and company.com is in the users search suffix, the correct vale for the NTLM host would be "**portal1**".

The following diagram shows NTLM based Captive Portal flow using a redirect. In the case of a transparent mode Captive Portal there would be no steps 2 or 5. Instead the firewall would spoof the

destination address and provide the 401 error code as if the target server had sent it.

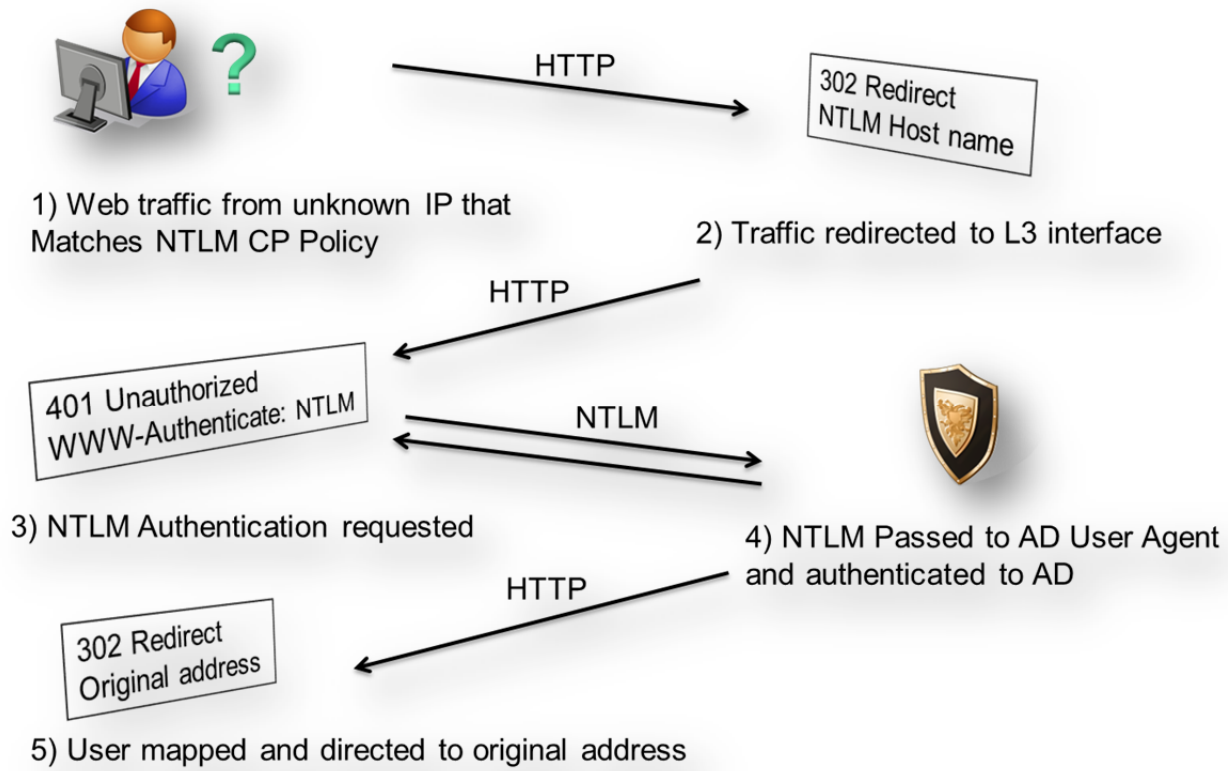

Web Form Captive Portal

This method displays a web page with fields for user name and password. The back end authentication can be RADIUS, LDAP, local database or native Kerberos. While this is the most disruptive user identification method it is also the method that will work with any kind of browser or operating system. As such it is an excellent method of last resort. The following diagram shows web form based Captive

Portal flow using a redirect.

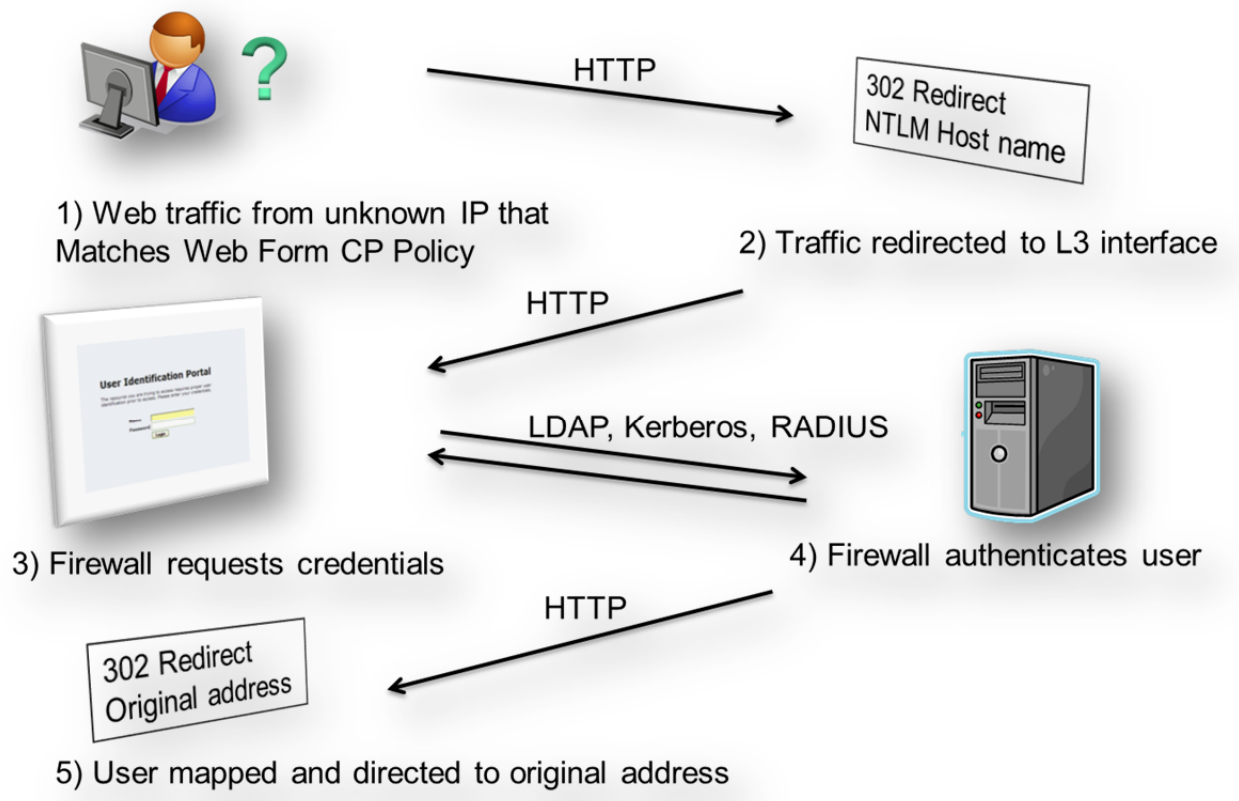

Certificate based Authentication

A user certificate can also be used by the captive portal to identify the user. Certificate based auth requires that trusted CA certs are loaded on the firewall and provisioned for user authentication. When the user first encounters captive portal they will be prompted for the certificate to pass on to the server. If no other authentication profiles are configured for the captive portal all further interaction between the browser and the portal should be transparent to the user. This is currently the only way to achieve fully transparent authentication for Linux and Mac clients using captive portal.

Best Practices for Configuring Captive Portal:

1.) Configure captive portal in redirect mode when possible. A single interface can be configured for L3 operations to host the portal for deployments using L2 or virtual wire.
2.) If using RADIUS insure the proper default domain is configured for users. If no domain is provided during the log in the default domain will be assumed.
3.) Kerberos authentication require less configuration for AD environments then LDAP and should be used in these cases.

4. Terminal Server Agent

The MS Terminal Server agent is a windows service that is installed on a Microsoft terminal server or Citrix server. The job of this agent it to intermediate the assignment of source ports to the various user processes. This source port information is passed on to the firewall and a user table is created including the user name, IP address of the terminal server and source ports of the users. This insures that each session from the terminal server is correctly mapped to the user that initiated it. No other user mapping features are required for these clients, although enumeration and group mapping still need to take place.

5. Palo Alto Networks client software

If the end point is running one of the Palo Alto Networks client software packages user identification will be provided by that software. There are currently 2 software packages that can run on the end point. Net Connect SSL VPN and Global Protect. Both of these packages will provide user information to the firewall they are connected to. No other method would be required to map the users to their IP addresses, though there would still need to be something in place to enumerate the users and their group membership.