



Simo Sorce  
Red Hat, Inc.

Red Hat Enterprise IPA, Technical Notes

# What is IPA ?

IPA stands for: Identity, Policy, Audit

IPA is about managing user/machine identities and related security policies

IPA v1 is focused around solving the problem of managing user identities across many machines on a network.

Based on a Red Hat Sponsored Open Source project:

- FreeIPA - <http://www.freeipa.org/>

# The Identity Management Problem

## Needs:

- Single point of Management (comprehensive view)
- Single source for Identities (duplication = mgmt hell)
- Single-Sign-On / Single-Password
- Single data store for auditing/reporting (compliance)

## Implementation problems:

- Integrated Management Interfaces
- Single point of failure
- Synchronization and/or Integration
- Distribution of data/credentials

# Attacking the Identity Problem

We need a storage mechanism that allows us to:

- Store identity information, in an extensible way
- Perform access control at the attribute level
- Organize Identities and allow group relationships
- Distribute Information across the enterprise
- Replicate Information on multiple servers



# Attacking the Identity Problem

We need an authentication method that allows us to:

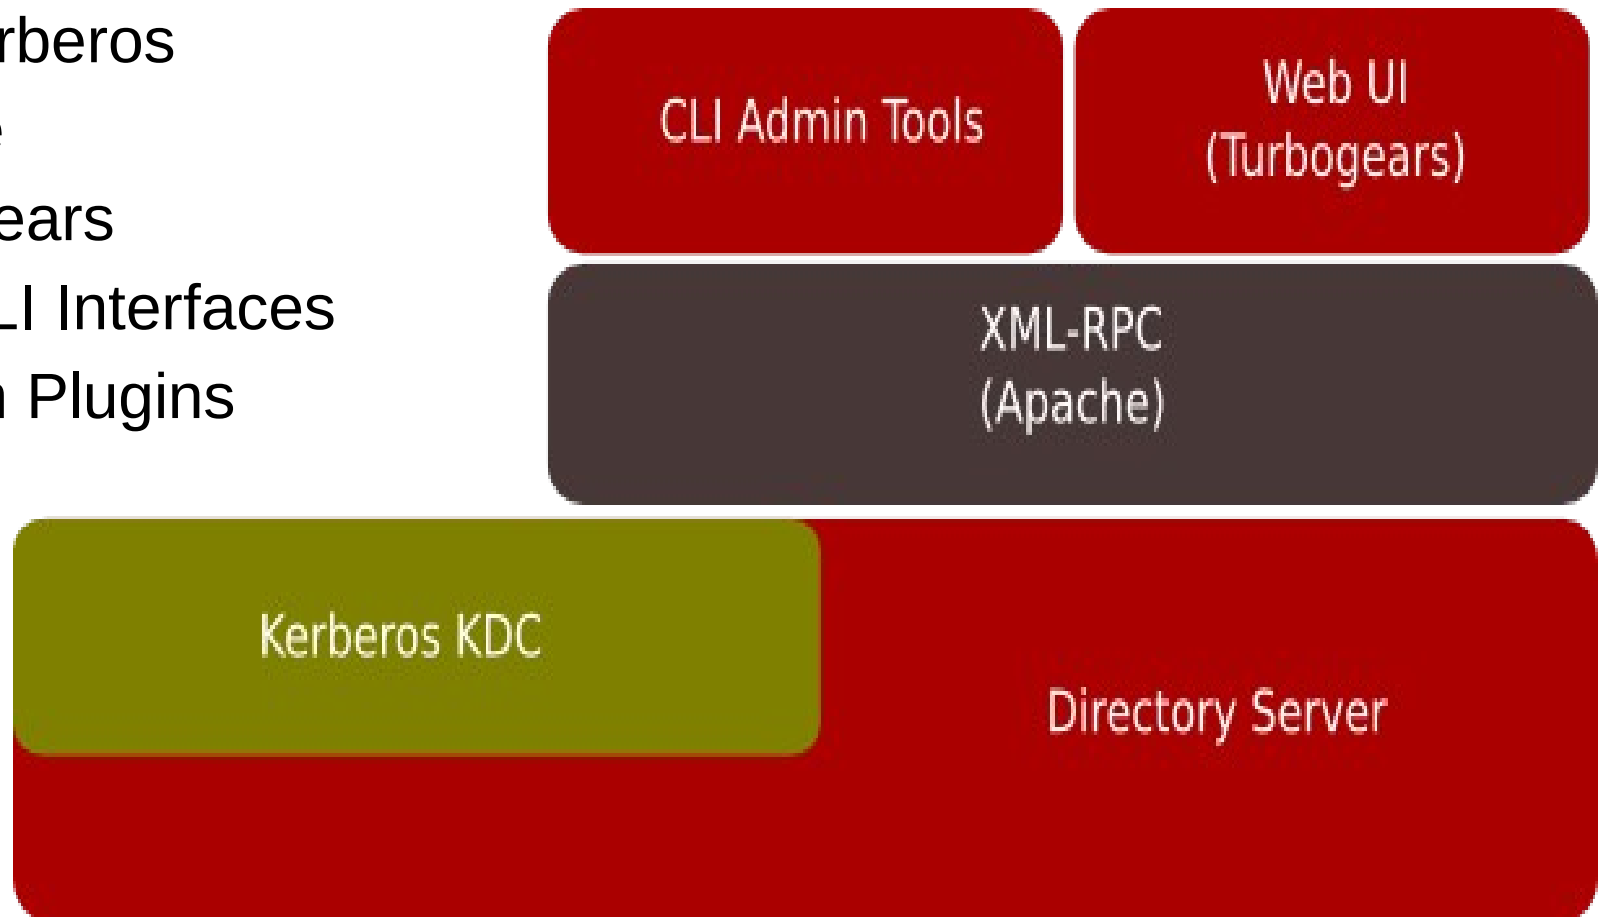
- Provide Single Sign On authentication
- In a way that allows administrators and users alike to carry on their identity while they access various services
- Is a tested Industry standard
- Is extensible/extended to use new authentication technologies like SmartCards



## Kerberos


# IPA Components

- Red Hat Directory Server
- MIT Kerberos
- Apache
- Turbogears
- Web/CLI Interfaces
- Custom Plugins
- NTP
- (DNS)



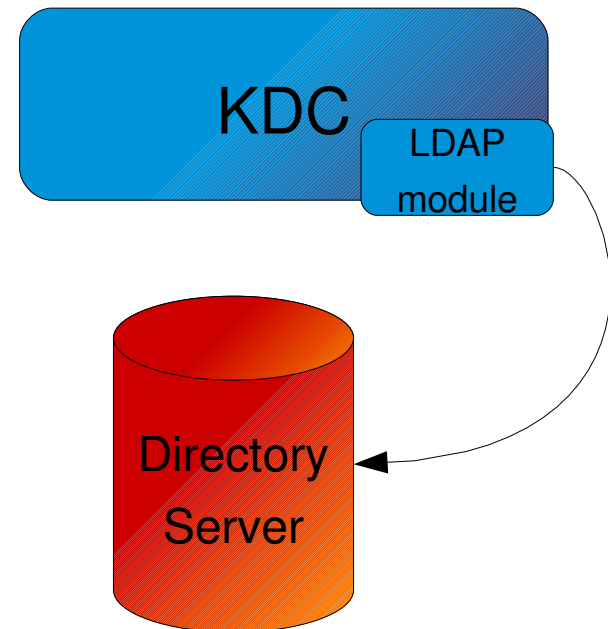
# Directory Server

- The Directory tree is split into 3 main containers
  - cn=kerberos
  - cn=accounts
  - cn=etc
- Under kerberos we store most of the kerberos related information, except the user principals
- Under accounts we store all our identities groups, services, and related information
- Etc stores configuration or system related data




# KDC

- The KDC is a standard MIT Kerberos Server as shipped with Red Hat Enterprise Linux 5, with the addition of an LDAP storage backend.
- The backend stores all kerberos principals, including expiration information and other principals related information. It also includes some realm related configuration
- Part of the KDC configuration is still saved in the classic kdc.conf file
- The kadmin daemon is deactivated as all management is done through the IPA management interfaces



# Management Interfaces

- User friendly Web Interface to create and manage users, groups and services.
- Command line tools with same functionality also available
  - ipa-useradd
  - ipa-usermod
  - ipa-groupadd
  - ...



The screenshot displays the Red Hat Enterprise IPA web interface. At the top, there is a red header with the text "RED HAT ENTERPRISE IPA". To the right of the header, there is a search bar with a dropdown menu set to "Users", a text input field containing "Type search terms here.", and a "Search" button. Below the header, a green bar indicates the user is logged in as "dkim".

The main content area is titled "Edit User" and includes a checkbox for "edit protected fields". Below this, there are two sections: "Identity Details" and "Account Details".

**Identity Details:**

- Title: Mr
- First Name: David
- Last Name: Kim
- Full Name: David Kim (with a "Remove" link and an "Add Full Name" link)
- Display Name: David Kim
- Initials: DK

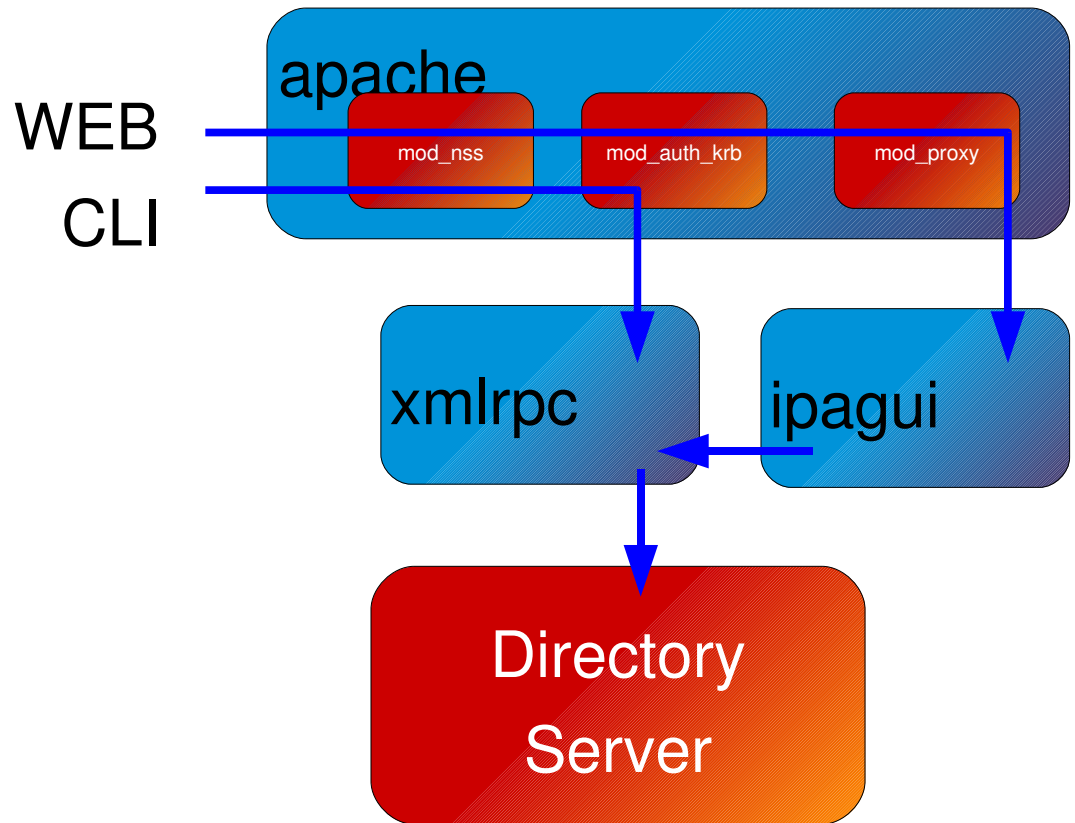
**Account Details:**

- Account Status: active
- Login: dkim
- Password: (empty field)
- Confirm Password: (empty field)

On the right side of the form, there is a "Tasks" sidebar with links for "Find Users", "Find Groups", and "Self Service". At the bottom of the form, there are "Cancel Edit" and "Update User" buttons.

# Management Interfaces

- Apache
  - mod\_nss
  - mod\_auth\_kerb
  - mod\_proxy
  - XML-RPC
- IPA-GUI
  - CherryPy
  - TurboGears



# Additional Components

- NTP

Time synchronization is extremely important, both Red Hat Directory Server and MIT Kerberos rely on clocks being synchronized on all machines joined to the IPA Realm, and especially on the IPA Masters.

Time synchronization is needed since the first installation to be able to create x509 certificates and other validity/expiration dates correctly.

It is needed for RHDS multi-master replication to work correctly.

It is need for Kerberos to work at all. Kerberos libraries and KDCs, by default, allow for a maximum of 5 minutes clock skew between machines for security reasons.

# Additional Components

- DNS

To allow easier configuration of the clients, DNS is used to provide basic information about the Kerberos Realm and the available IPA Servers using standard TXT and SRV records.

To be able to retrieve Tickets to access remote machines, clients need to be able to resolve machine names.

Because some services might provide only an IP address to the kerberos libraries, Forward as well as Inverse resolution should work.

A DNS Server is not integrated in IPA v1, but a correctly configured DNS infrastructure is required.

# Server configuration

## Very simple configuration

- Check the machine is ready (hostname, IP, base OS)
- Install rpms
- Run ipa-server-install
  - Answer a few questions:
    - DNS Domain and Realm name (defaults suggested)
    - Directory Manager password (required)
    - Admin User Password (required)
    - Done!
- The installation program configures all necessary components: NTP, Directory Server, Kerberos, apache, ipa-kpasswd, ipa-gui, client side bits

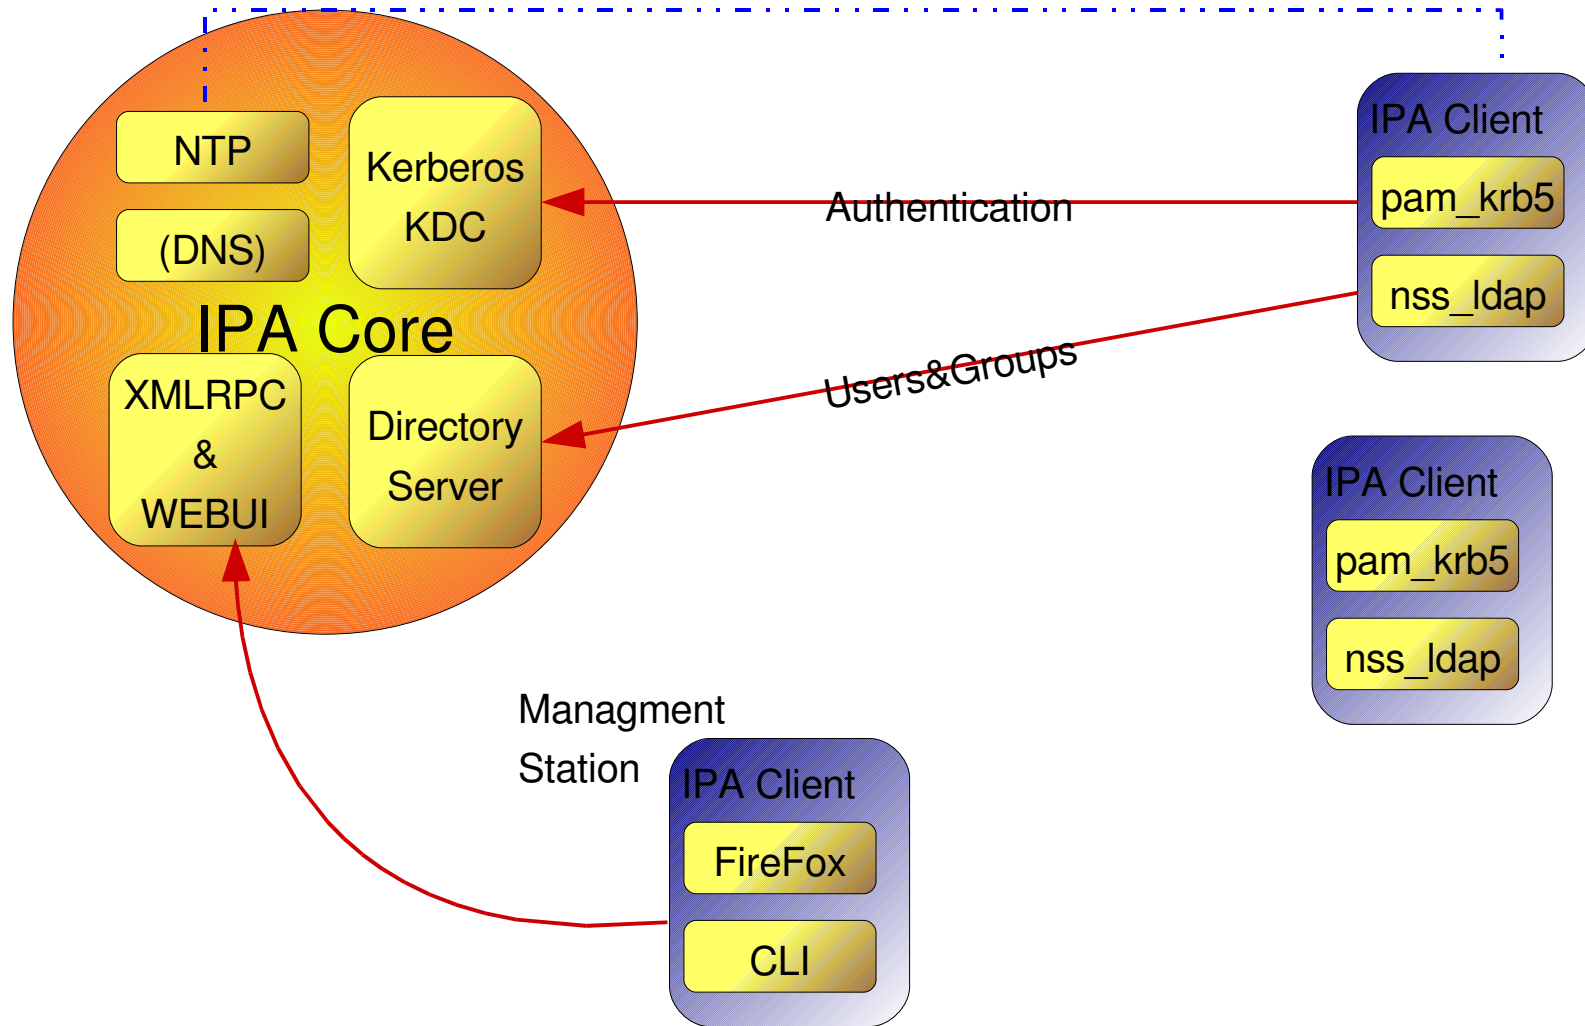
# Client configuration

We provide configuration scripts for Fedora and RHEL4/5 clients, for other platforms, configuration instructions are provided

The ipa-client-install script is very simple

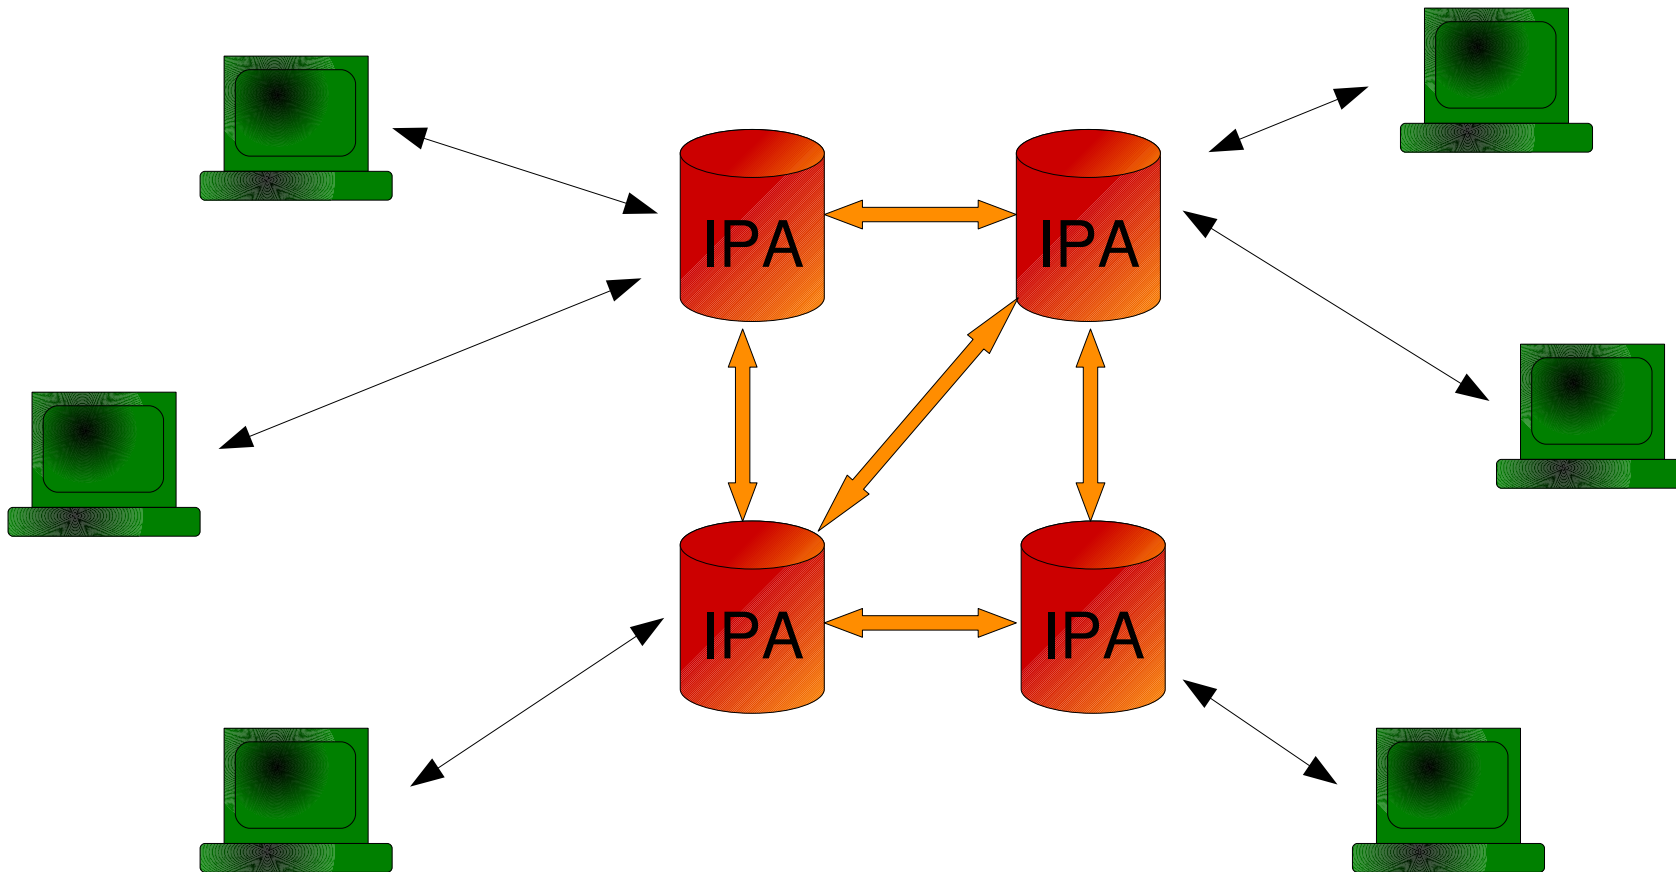
- Discovers IPA Server using DNS SRV records, falls back to request IPA Server name in case of failure
- Configures subsystems accordingly
- ntp.conf
- krb5.conf
- ldap.conf
- Pam stack
- nsswitch.conf (nscd)

# IPA v1 network diagram



# IPA v1 network topologies

Up to 4 masters, replicating informations to each other



# Moving Forward

So far in v1 we addressed only the User Identity problem and we provide a very simple interface to manually create service/machine principals for use on client and servers.

In v2 we plan to complete the Identity functionality by managing machine identities. We are also planning on starting implementing the other 2 letters of 'IPA', Policy and Audit

# Machine Identity

In a kerberos infrastructure, machines as well users and services need identities.

What is a machine Identity needed for ?

- User login stronger credentials validation
- Single-Sign-On SSH login using GSSAPI and Kerberos tickets
- Authenticated, Signed and Sealed access to the directory to retrieve user information

# Machine Identity

Machines need to obtain a kerberos keytab containing a principal named after the machine's Fully Qualified DNS Name.

Ex: [host/ipa-client-1.example.com@EXAMPLE.COM](#)

This requires 2 things:

- Each machine need a persistent DNS name resolvable by other client in the network
- Each machine needs to request a machine keytab to the IPA Server

# Machine Identity and DNS

A permanent DNS name is required, this means that administrators need to keep track of all their machines and make sure the DNS name correctly reflects the IP address (and vice versa). Failing to keep the DNS updated would prevent the kerberos mechanisms from functioning correctly.

We plan on integrating a DNS server in the next versions of IPA, like we did for the KDC and the Directory Server to make it easier to correctly keep track of machines through a unified data storage and management interface

# Client Enhancements, Policies and Audit

We currently use `pam_krb5` and `nss_ldap` on the clients to perform authentication and retrieve users and groups information. While these components work well enough in general they also present well known deficiencies:

- Lack of server affinity
- Cannot operate offline
- Inefficient caching and resource usage

We plan on building a new client service to be installed on the clients that will provide a solution to the mentioned problems and a few compelling features:

- Machine credentials management, offline caching, server affinity, policies, central auditing

Questions ?

Thank you!