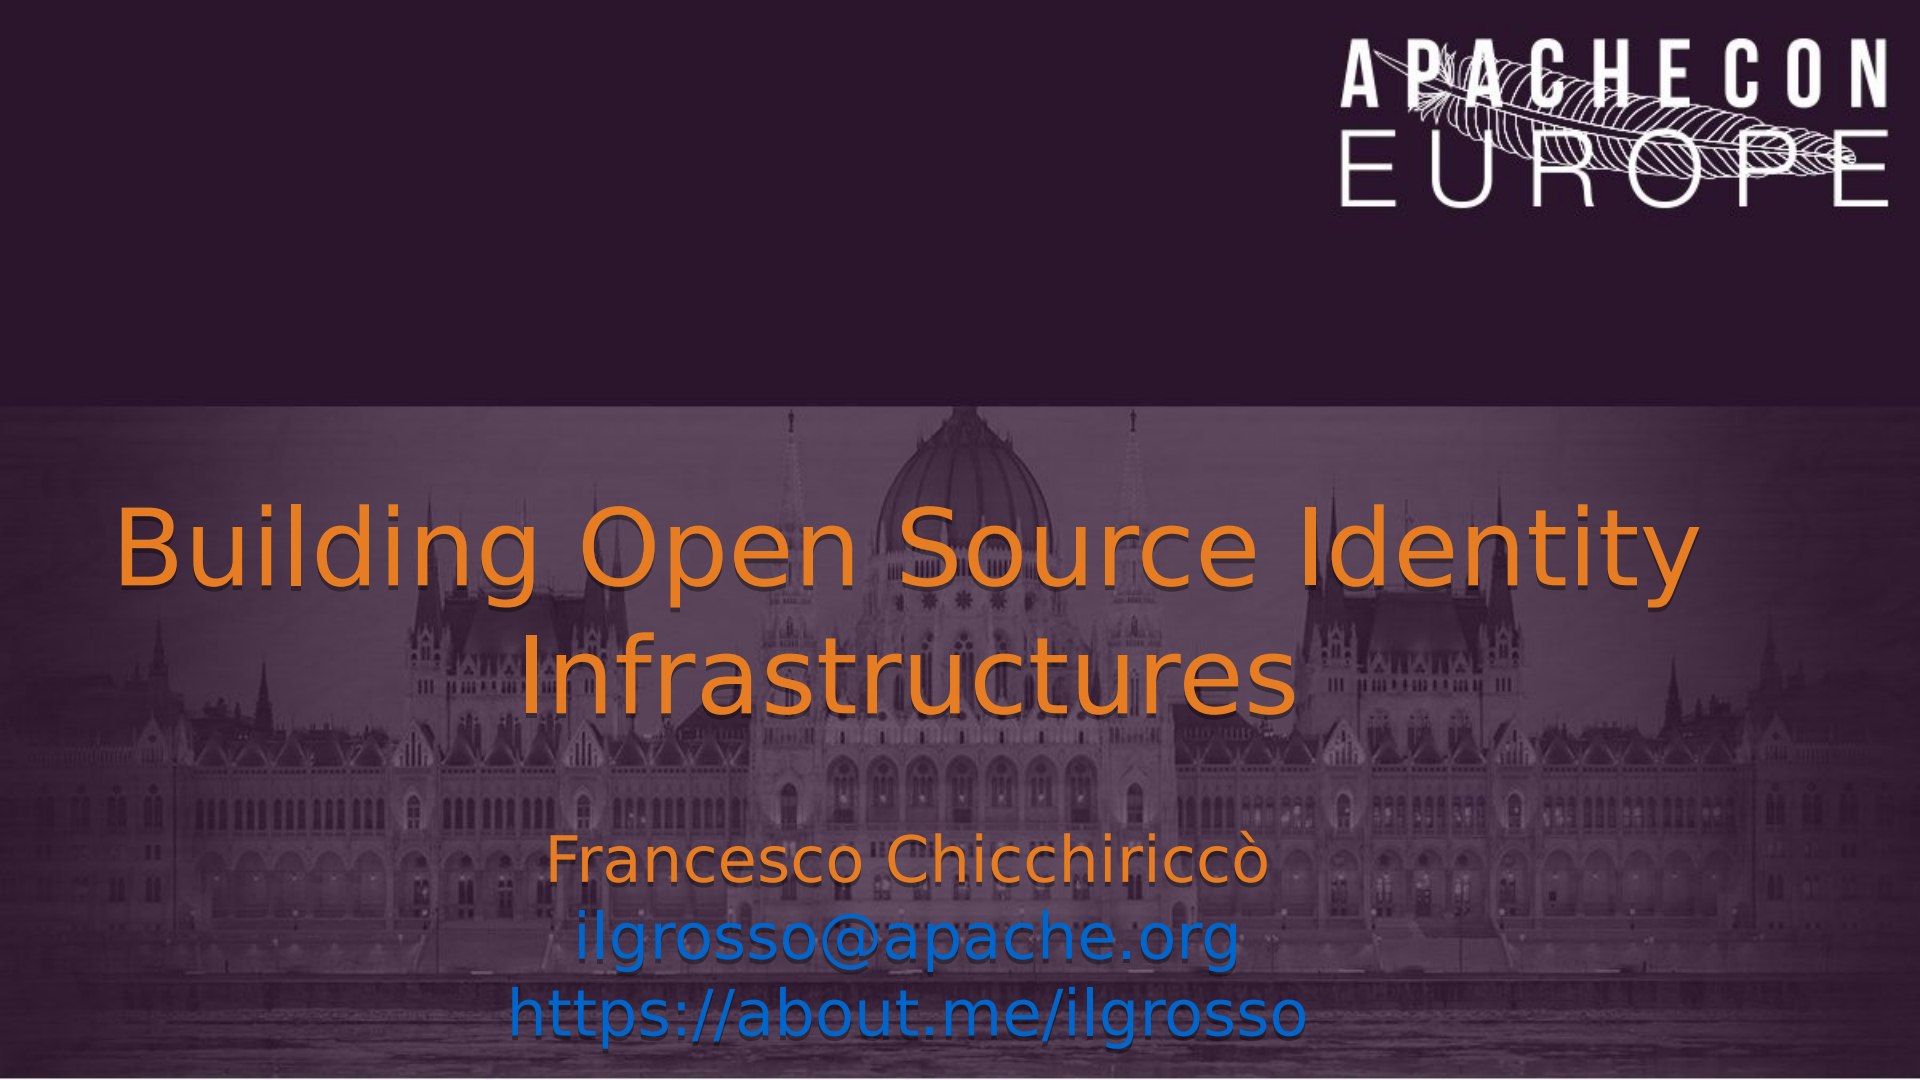

Building Open Source Identity Infrastructures

Francesco Chicchiriccò
ilgrosso@apache.org
https://about.me/ilgrosso

The Identity Management Need

Identity Vs Account

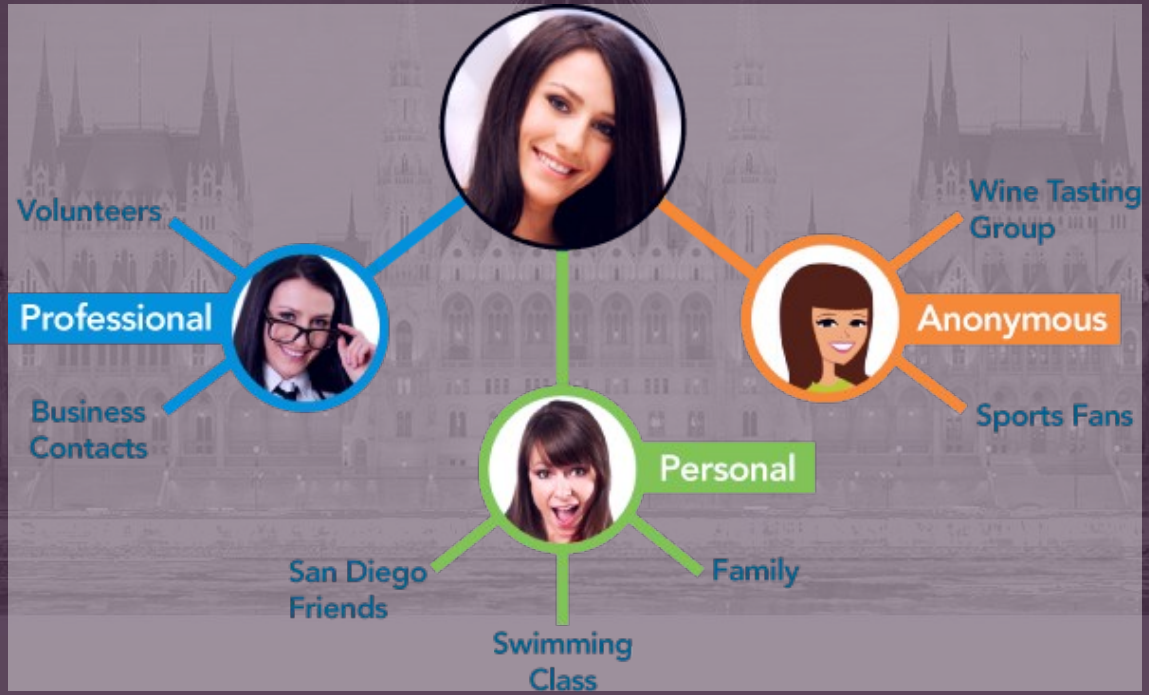

Source: https://saberhamidi.wordpress.com/2015/02/22/topic-2-should-we-have-more-than-one-online-identity/

Identity Vs Account

- **Account**
  - record containing data about a person
  - technical info needed by the information system for which the account is created and managed
- **(Digital) Identity**
  - representation of a set of claims made by one digital subject about itself
  - ...it's you

Why Identity Management?

- Operational costs
  - Multiple sources of identity data
  - Manual user provisioning and password reset
  - Labor-intensive, paper-based approval
- Compliance
  - No record of who has access to which IT resources
  - Difficult to deprovision access rights upon termination
  - No complete audit trail available
  - Hard to prevent unauthorized access

Which identity?

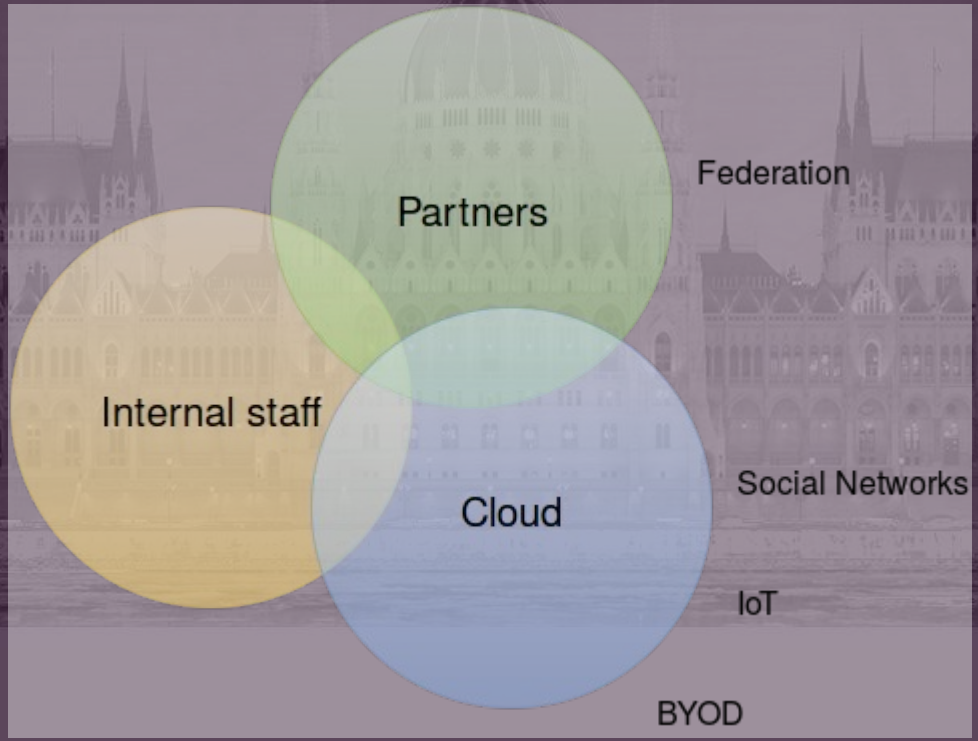

Identity Solutions

Identity Technologies

- Identity Stores
  - Storage of user information
- Provisioning Engines
  - Synchronize account data across identity stores and a broad range of data formats, models, meanings and purposes
- Access Managers
  - Security mechanisms that take place when a user is accessing a specific system or functionality

Identity Store

- Examples
  - LDAP / Active Directory
  - RDBMS
  - Meta and Virtual Directories
- Accounts can be created and managed in one place only
- Each application manages authentication separately
  - The user may use the same password for all the connected applications

...is it enough?

- Heterogeneity of systems
- Lack of a single source of information
  - HR for corporate id, Groupware for mail address, ...
- Need for a local user database
- Inconsistent policies
- Lack of workflow management
- Hidden infra management cost, growing with organization

Provisioning Engine

- Keeping the identity stores as much synchronized as possible (and practical)
- Need to be customizable and flexible
- Priority: non-intrusive
- Focused on application back-end
- Critical: data exchange with identity stores
  - Connectors
  - Agents

Identity Lifecycle

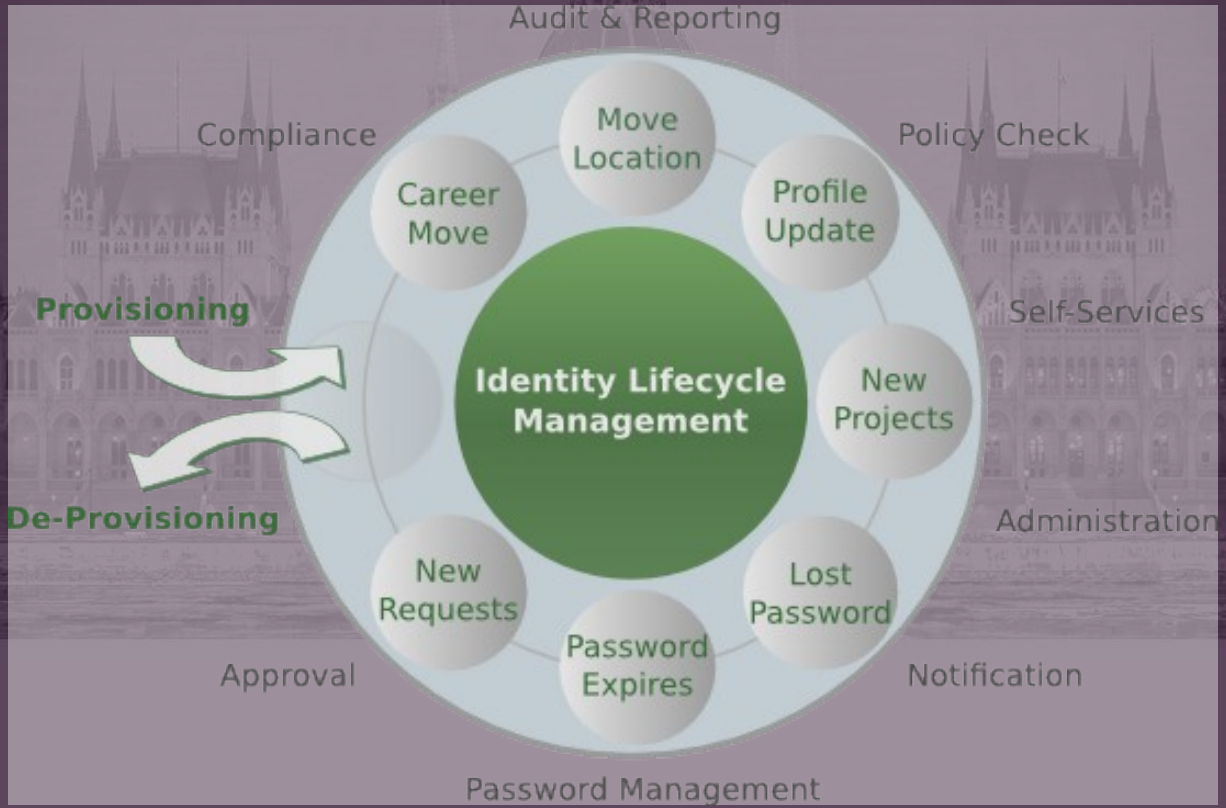

Access Manager

- Mediator to all access to all applications
- Focused on application front-end
- Aspects
  - Authentication
    - Single SignOn
  - Authorization (OAuth, XACML, ...)
  - Federation (SAML, Liberty, ...)
- Mainly applicable to web applications
- Difficult integration with pre-existing apps

Reference Identity Scenario

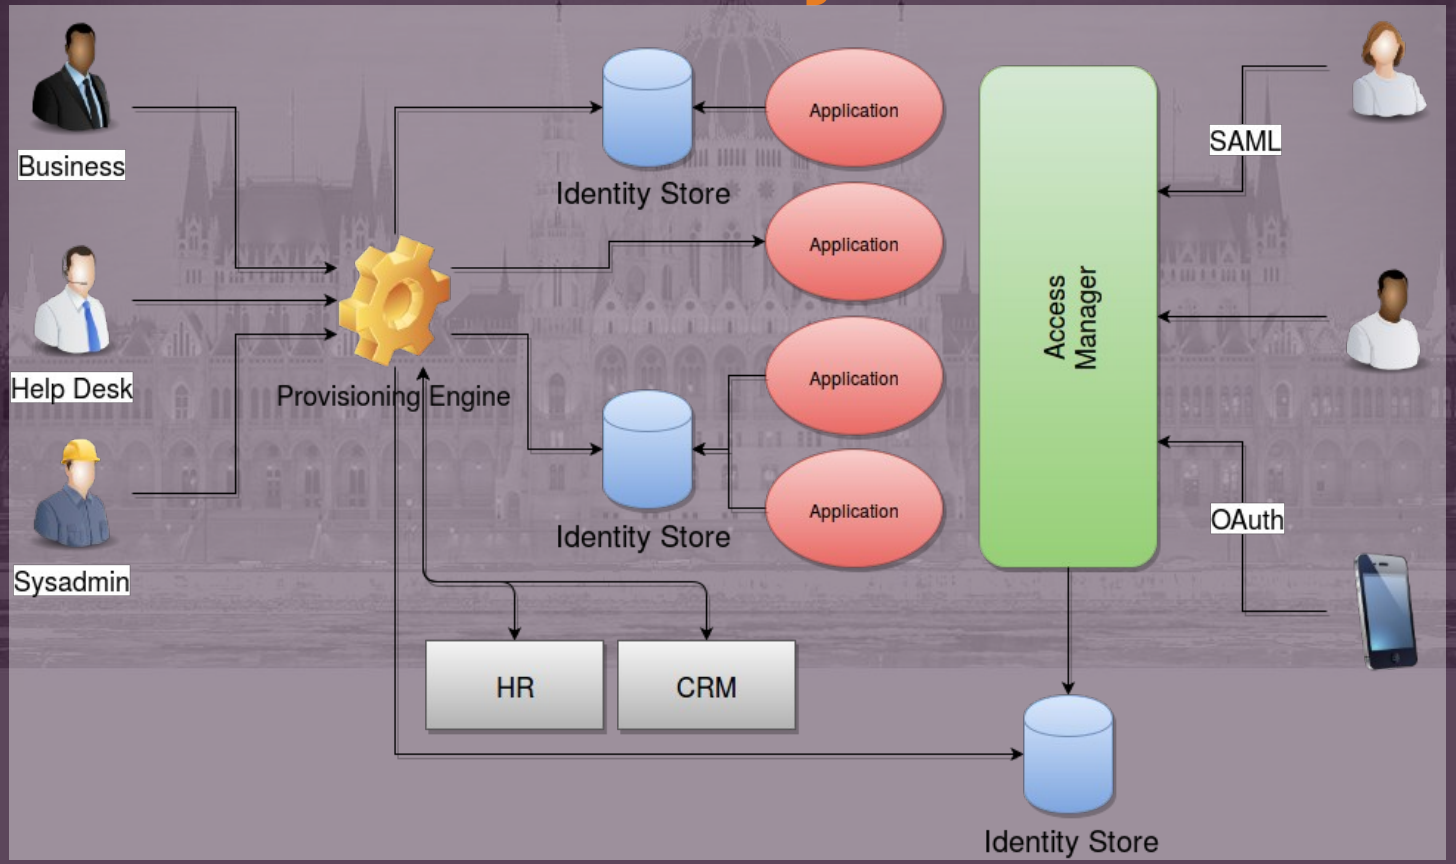

Identity Infrastructures

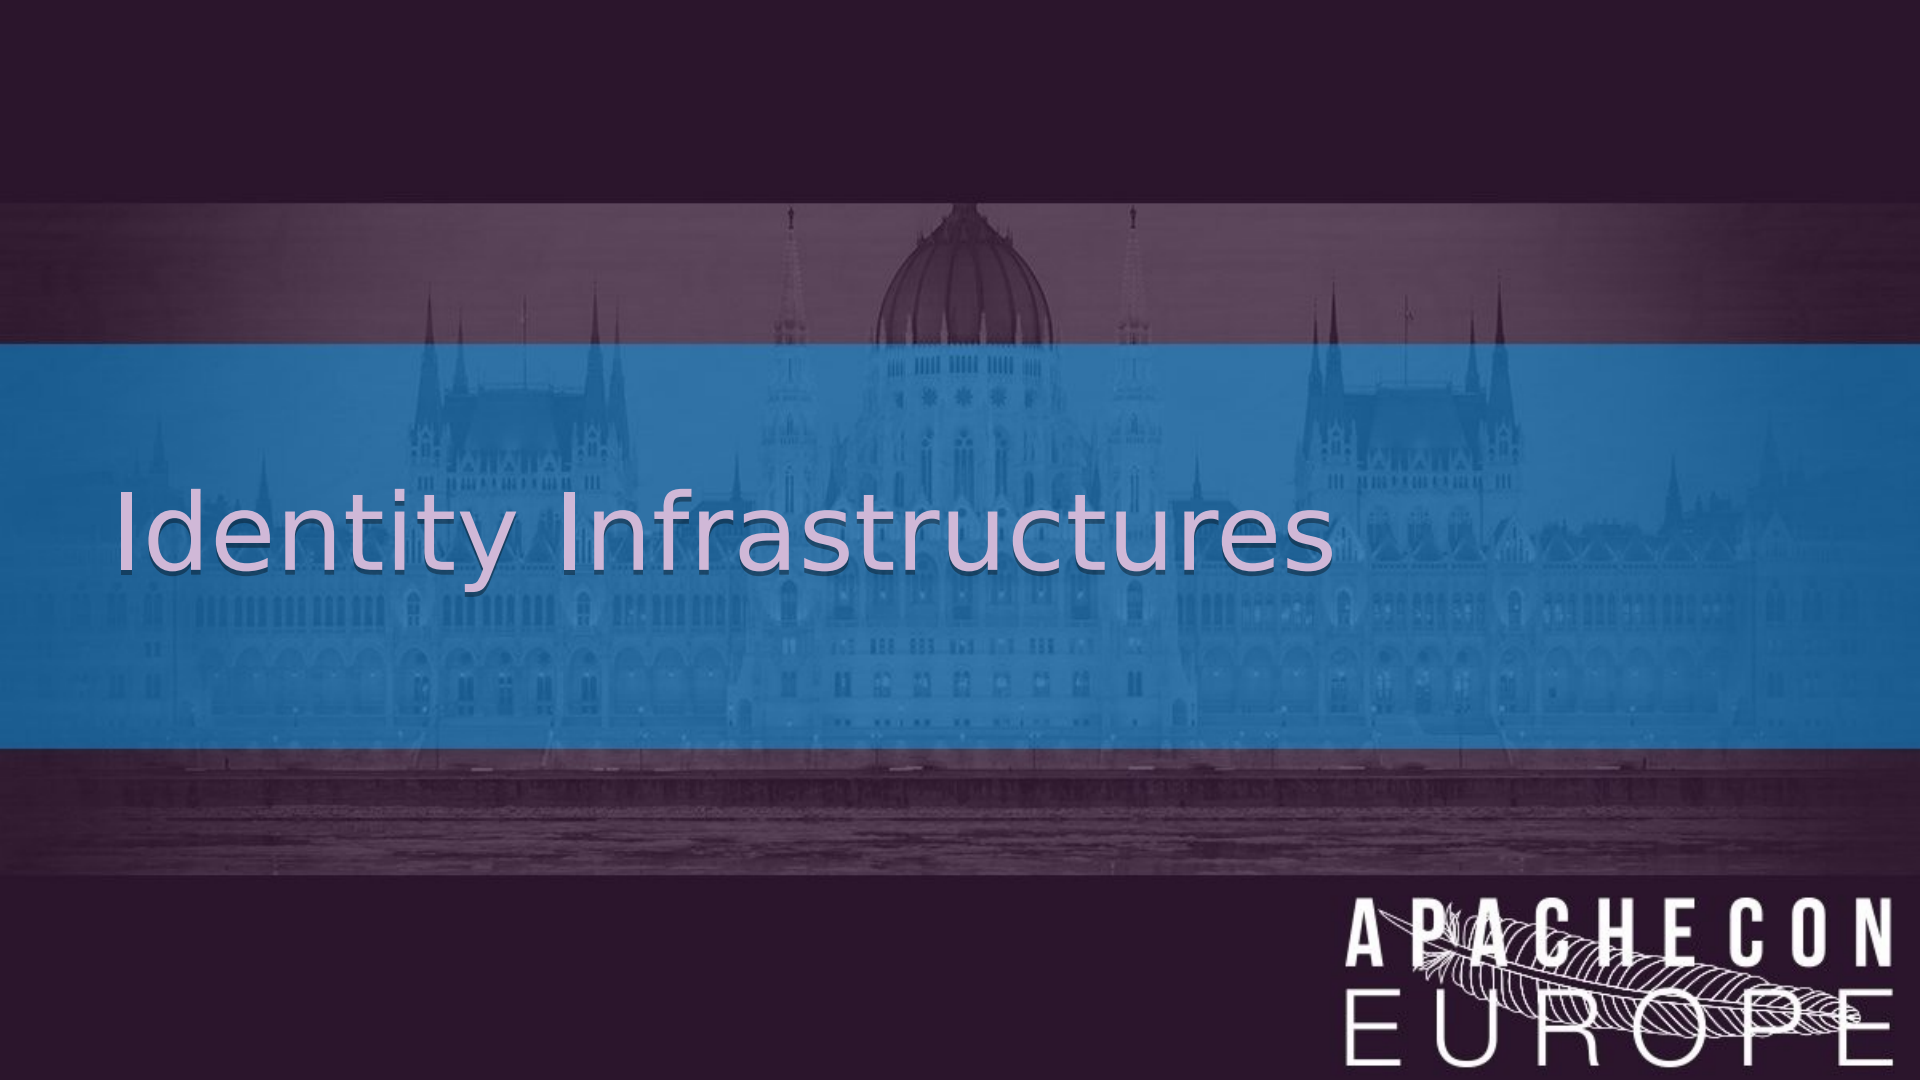

Gather...

- Number and type of identities
- Number of roles / groups (and what are they used for)
- External resources (all covered by standard connectors?)
- Approval workflow(s)?
- Self-service?
- Which applications to protect?
- Which authentication mechanisms?
- Which authorization types?

...essentially, shape the identity and access flows

...design...

- Schema for various identities (users, roles, groups, ...)
  - Identify mapping for all resources
  - Not too complex!
- Watch roles size to avoid RBAC's role explosion
- Don't be tempted to redesign the whole network
  - Provisioning needs to be flexible
  - Reduce impact of access management on existing applications
- Prioritize requirements

...build...

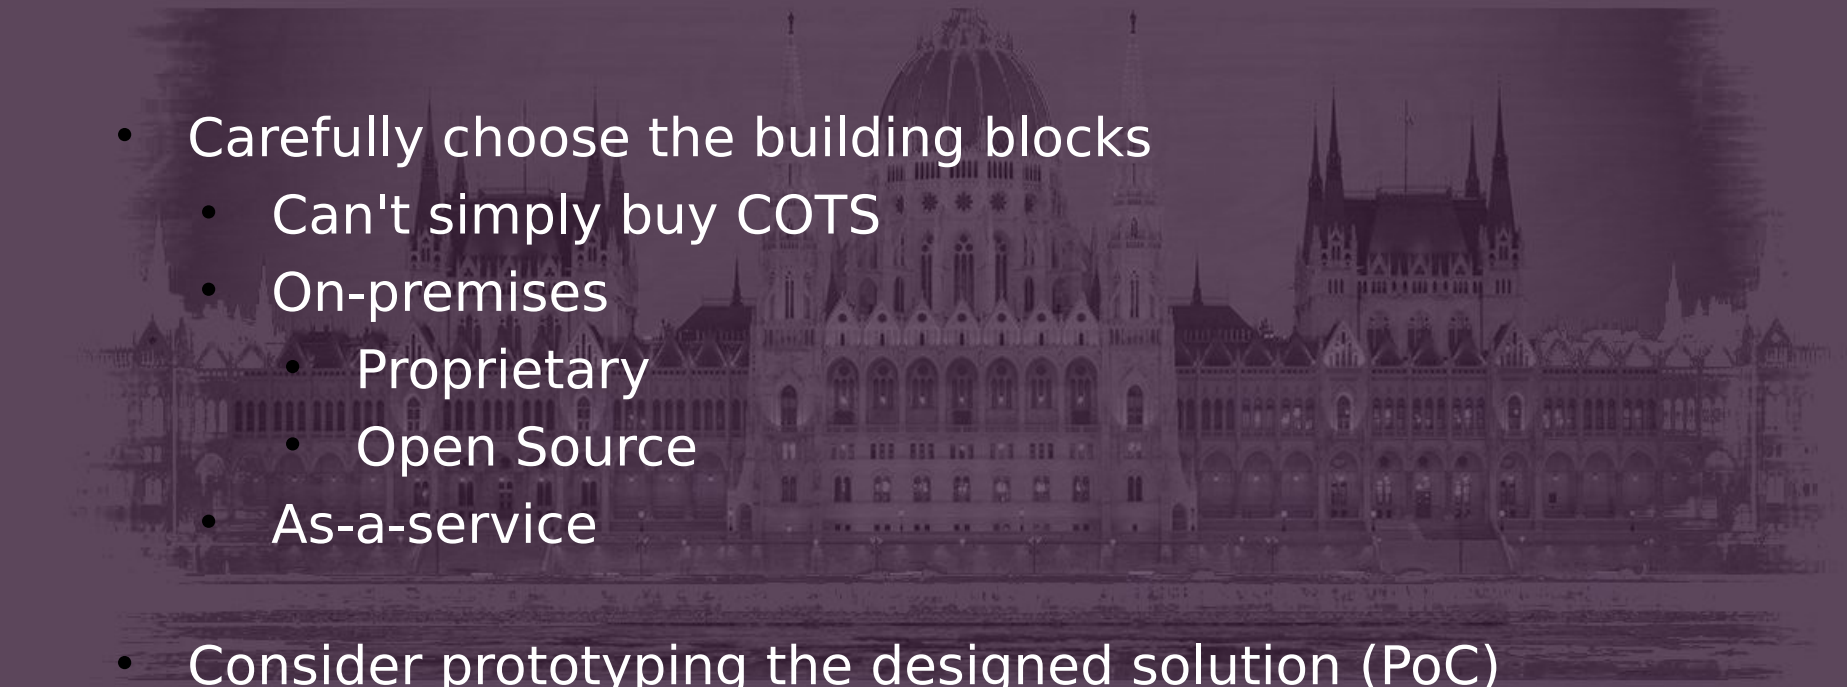

- Carefully choose the building blocks
  - Can't simply buy COTS
  - On-premises
    - Proprietary
    - Open Source
  - As-a-service
- Consider prototyping the designed solution (PoC)

...and start again

- IAM is a continuous process, not a turn-key project
  - New applications to protect
  - New resources to integrate
  - Identity flows evolution
- IAM deliveries frequently fail
  - Mix of complex and unrelated technologies
  - Unexpected interactions
  - Mess with internal processes
  - Discover Policy Vs Reality

The Open Source Identity Stack

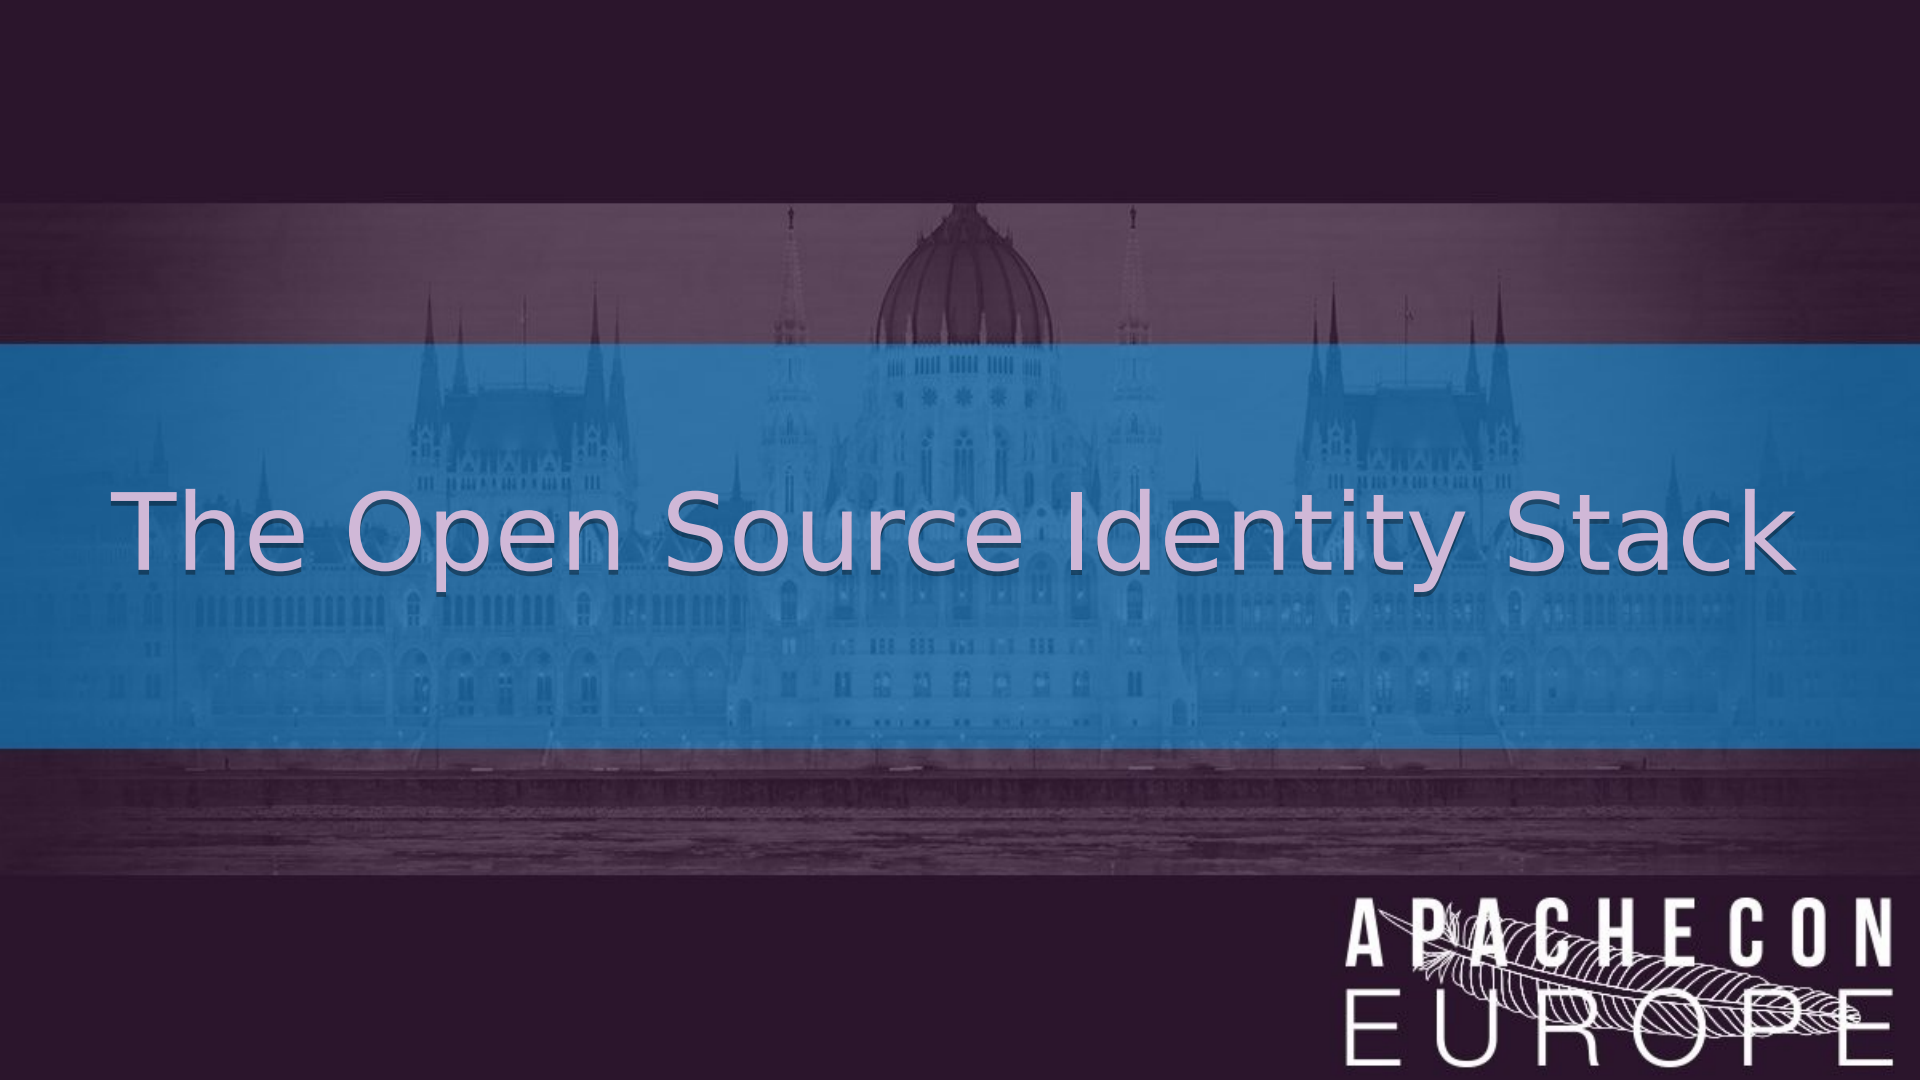

Open Source IAM

- Why?
  - Flexibility, adaptability and agility
  - Cost effectiveness
  - Start small and grow
  - Solid information security
  - No vendor lock-in
- Caveats
  - Integration with proprietary software (AD over all)
  - Enterprise support availability

Available Components

Selection Criteria

- Open Standards
- Design for integration
- Well-established
- Supported
- Alive
- ...Open Source!

The Identity Ecosystem

- Triggered by open companies in the Open Source IAM area
- Common place for open source players, system integrators and service providers
- Ensuring IAM open source components work well together
- Easy access to enterprise support providers
- Several options for each single component

- More at http://www.identity-ecosystem.org/

Real World Use Cases

Disclaimer

I am V.P. Apache Syncope and CEO of Tirasa, providing enterprise support and services for Apache Syncope, so...

don't be surprised Syncope is everywhere :-)

#1 Stadtwerke München

- One of largest German municipal utilities
- Mobile ticketing for public transportation and bike sharing
  - self-registration
  - login
  - password reset
  - user suspend / reactivate
- > 250k registered users
- > 80k authentications per day

#2 Ospedali Riuniti Ancona

- University hospital
- Active synchronization from HR to Microsoft Active Directory
- Centralized provisioning, authentication and authorization of medical record systems
- Windows domain SSO
- SAML 2.0 federation with regional network
- ~ 5000 users

#3 Stichting Bibliotheek.nl

- Dutch foundation that aims to expand and manage the Digital National Library
- The IAM infrastructure aims to hold all users of the national library in the Netherlands, fed by a continuous feed from the local libraries
- All Dutch library members can authenticate and use digital services connected to the IAM infrastructure
- > 8 million users

#4 University of Milan

- Very complex provisioning flows involving
  - Microsoft Active Directory
  - OpenLDAP
  - 3 different RDBMS
  - Oracle E-Mail Server
- ~ 5k employees
- > 60k students
- ~ 800 roles

Questions?