Penn State University Shibboleth Experience

Webassign & Napster

Mark Earnest
Lead Systems Programmer
Academic Services & Emerging Technologies
Pennsylvania State University
mxe20@psu.edu

PENNSTATE

In the Beginning…

- What we had:
  - DCE based central authentication
    - ~184,000 Principals
  - IBM LDAP Directory (SecureWay)
- Early Shibboleth Experience
  - Beginnings around spring '02
  - Renee Shuey and I were tasked by bringing up a test Shibboleth Origin
  - Shibboleth Version: Alpha 2.5
  - Origin site brought up Summer '02
    - Ran from my desktop, 800MHz P3 – 256MB RAM

PENNSTATE

Webassign

- Web based physics resource
  - Hosted at North Carolina State University
- Prior to Shibboleth, separate accounts were maintained at NC State
  - Obvious problems
- Summer '02 a limited Shib trial occurred with ~20 students
- Fall '02 a more dedicated trial was launched with one Physics class (~200 Students – 3 sections)
  - Shibboleth Beta 1 (Still from my desktop)
- Spring '03 Shibboleth was rolled out to all Webassign using Physics classes (~1800 Students)
  - Shibboleth 0.7 (Finally on a real machine)

PENNSTATE

Webassign Details

- Two attributes passed:
  - eduPersonPrincipalName
  - eduPersonAffiliation
- Accounts still created manually
  - Ideally we wanted dynamic account creation
  - This required more attributes that we had in LDAP
- No scaling or performance problems
- Tomcat proved to be a bit unstable at that time
  - Occasionally stopped responding and needed to be bounced

PENNSTATE

Webassign Phase II

- Spring '03 – Dynamic account creation
- New LDAP Attributes
  - eduPersonEntitlement (course list)
    - URN:PSU.EDU:COURSE:UP:PHY002:001
      - LONG list – Regexps for attributes in Shib 1.2
  - CN (full name)
- Accounts created based on course enrollment
- Production Origin site now run by AIT
  - IBM HS20 Blade – 2.4GHz – 2.5GB RAM
- Shibboleth 1.1

PENNSTATE

Napster

- PSU Entered into an agreement with Napster to provide music service to students
- Shibboleth quickly appeared as obvious solution to registration problem
- New Problems:
  - Access limited to Residence hall students during this trial
  - Convey identity while preserving privacy
  - SCALE!
- PSU and Napster agreed early on that identifiable data should not be sent, but they still needed to some form of persistent identifier.

PENNSTATE

Napster Attributes

- New attributes were used for Napster
  - eduPersonEntitlement – URN:PSU.EDU:MUSIC
    - URN:PSU.EDU:SHIBFIX – null attribute problem
      - Fixed in 1.2
  - eduPersonTargetedID – Opaque Handle
    - Hash of principalname, target name, and secret seed value
    - We generate these on the fly, ideal way would be to generate them at once and store them in LDAP
- Attributes we DID NOT use:
  - eduPersonPrincipalname, CN, etc.
- The entitlement attribute was populated based on our “HOTL” code for residence halls
  - Problematic, ran into some inaccurate data

PENNSTATE

Load Testing

- To prove to ourselves we could handle this, we did some very unscientific load testing
  - PERL script – many fired off at once and timed
  - Response goal was under 5 seconds
    - Maximum of 25 simultaneous connections
- 500 – Internal Server Error
  - We originally did not check for this, later found out we should have!
  - Setting the acceptCount on Tomcat lower than MaxClients on Apache is a bad idea…
    - Tomcat will not queue, returns errors when overwhelmed

PENNSTATE

Load Balancing

- In order to ensure success, we relied on load balancing
  - 25 simultaneous connections likely more than enough, but we really needed to be sure
  - Redundancy was also an issue
- 5 Blades dedicated to Napster Origin
  - OVERKILL!
  - Load balancing performed by Cisco's SLB
- Persistent Opaque Handle
  - Required because HS & AA requests likely used different blades
  - Principalname encrypted with secret seed value

PENNSTATE

Questions?

- Questions?
- Thank you :)