LEVERAGING THE CLOUD FOR NIST SP 800-171

Preston Smith
Director of Research
Computing Services

PURDUE UNIVERSITY

May, 2017 – HTCondor Week, 2017

OVERVIEW

- Background
- Motivation
- The Cloud
- REED Overview
- Future Wrok

BACKGROUND

2006!

- Director of Research Computing Services
- I was once the “Condor Guy” from Purdue –
  - (I did my thesis on ROI of campus Condor grid)
- Now I lead the group supporting computing for scientists all over campus

- **Disclaimer:** I’m not the security/compliance guy!

PURDUE
UNIVERSITY®

HTCONDOR AT PURDUE

- Purdue has been an HTCondor site for a long time
  - Backfill scheduler on Community Clusters
  - Teragrid Condor resource
  - Large Campus Grid
- Today, HTCondor is used at Purdue in more specialized use cases….

TRENDS AT PURDUE

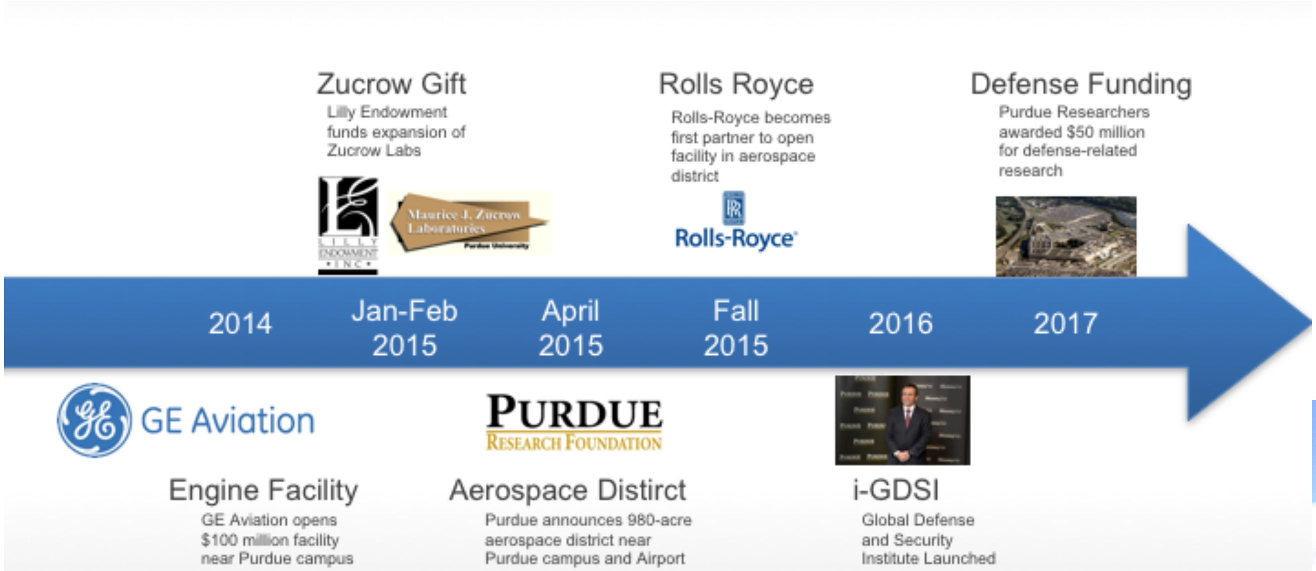

Zucrow Gift
Lilly Endowment funds expansion of Zucrow Labs

Rolls Royce
Rolls-Royce becomes first partner to open facility in aerospace district

Defense Funding
Purdue Researchers awarded $50 million for defense-related research

2014 | Jan-Feb 2015 | April 2015 | Fall 2015 | 2016 | 2017

Engine Facility
GE Aviation opens $100 million facility near Purdue campus

Aerospace Distirct
Purdue announces 980-acre aerospace district near Purdue campus and Airport

i-GDSI
Global Defense and Security Institute Launched

http://img.scout.com/sites/default/files/2016/05/18/4a2dccb34e894e0c3cd6f8f58862d32c.jpg

ITAR COMPUTING

- Since 2010, Purdue has offered a small cluster for ITAR contracts
    - Limited physical locations
    - “US persons” only
    - Also ITAR data storage
- Largest user group runs high-throughput Condor workflow with DAGMan managing

ITAR ++

- Researcher using ITAR system received new requirement for compliance with
  - DFAR 252.204-7012 (b)(2)(ii)(D) *
  - Sponsor wouldn’t negotiate on the requirement

* Guidance to Stakeholders for Implementing Defense Federal Acquisition Regulation Supplement Clause 252.204-7012 (Safeguarding Unclassified Controlled Technical Information)

CREATING A SEPARATE ENVIRONMENT

"If nonfederal organizations entrusted with protecting CUI designate specific information systems or system components for the processing, storage, or transmission of CUI, then the organizations may limit the scope of the CUI security requirements to those particular systems or components."

(NIST SP 800-171, pg. 3)

COMPLIANCE WITH DFAR-7012

- Was… problematic
    - All of campus?
    - Datacenter constraints to bring ITAR system in line
    - Certification/Audit for compliance prohibitive

Maybe the cloud?

NOT JUST ANY CLOUD

- If using Cloud, it must be equivalent to FedRAMP-Moderate.
    - 325 controls, selected from NIST 800-53
    - Accredited by approved 3rd party

WHY GOVCLOUD?

- GovCloud required for CUI data? **<u>NO!</u>**
  - Most Purdue contracts with CUI also include ITAR provisions – US persons
- Amazon discourages GovCloud for CUI unless there is an ITAR requirement
- US data centers, maintained by US persons only
- Fewer resources, fewer tools, higher cost

CLOUD PROVIDERS ARE NOT TOTALLY TURN-KEY

- Many carry certifications of some kind
  - *ISO, PCI, HIPAA, FISMA*
- This applies only to base infrastructure and/or managed services
- **YOU** are still responsible for ensuring compliance of what you build inside.

PURDUE'S SOLUTION

REED

- **Research**
- **Environment (for)**
- **Encumbered**
- **Data**

- A hardened environment in AWS GovCloud built to NIST SP 800-171 that allows for processing of data with strict regulations

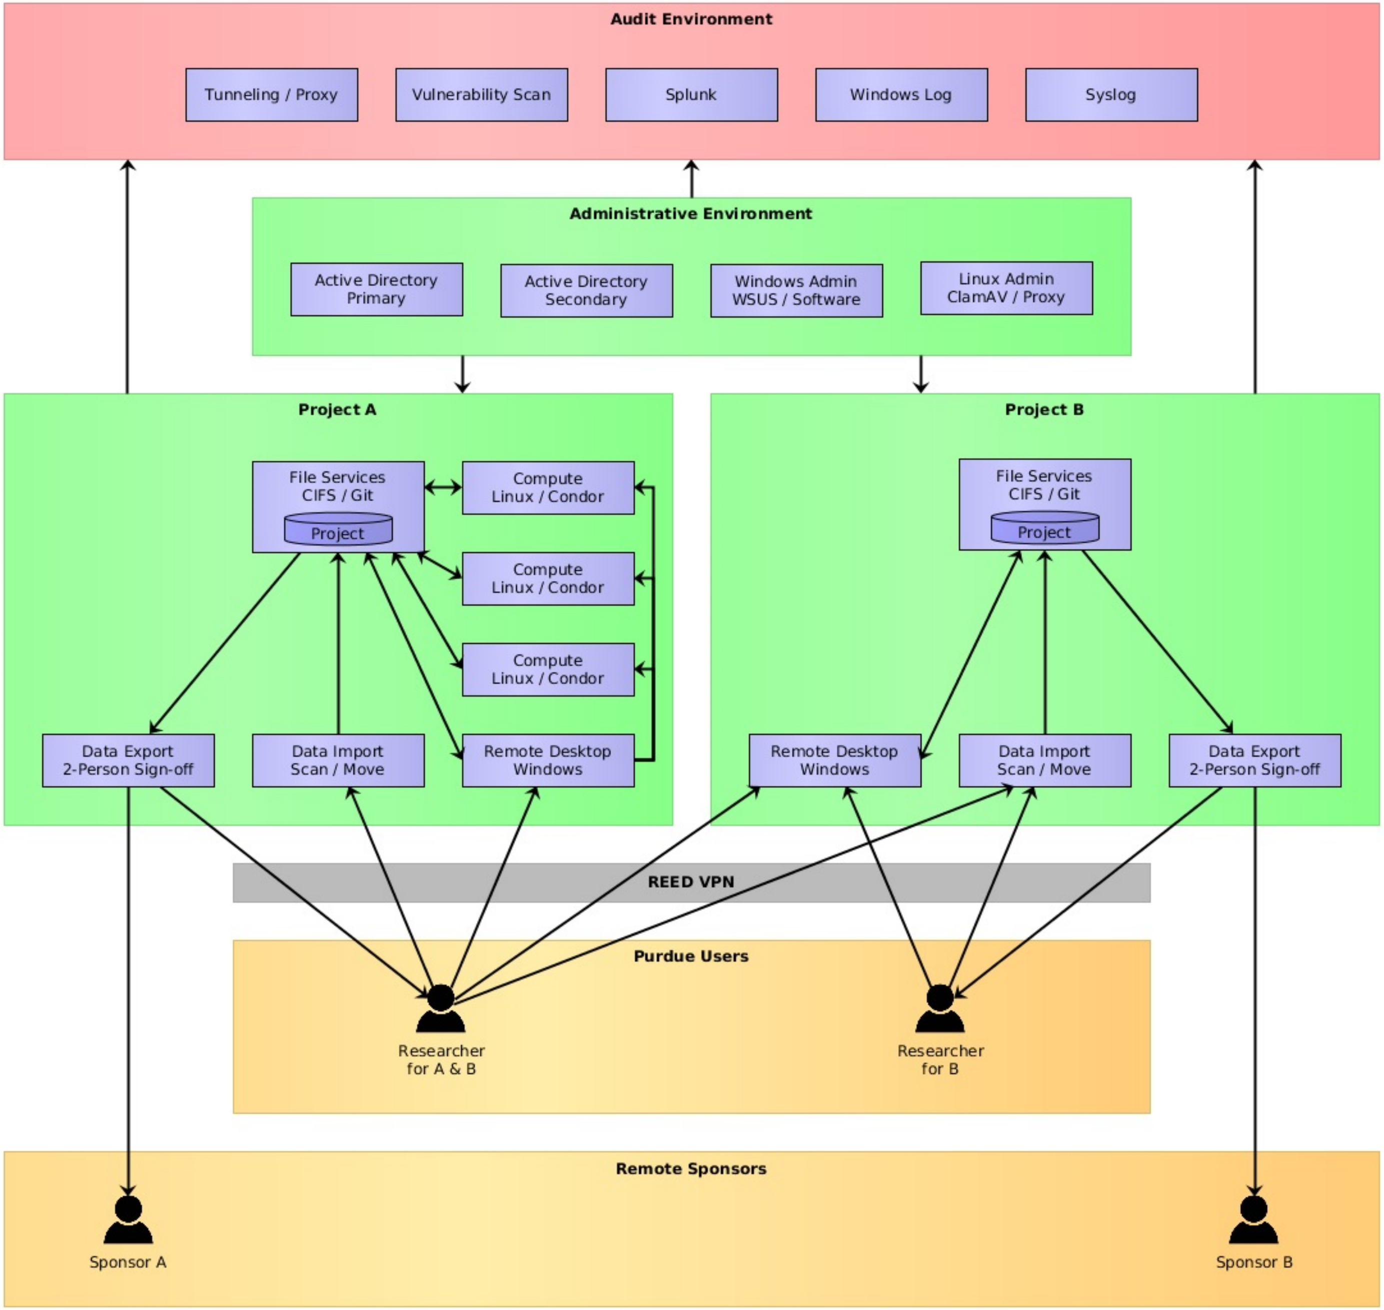
Audit Environment
Tunneling / Proxy
Vulnerability Scan
Splunk
Windows Log
Syslog
Administrative Environment
Active Directory Primary
Active Directory Secondary
Windows Admin WSUS / Software
Linux Admin ClamAV / Proxy
Project A
File Services CIFS / Git
Project
Compute Linux / Condor
Compute Linux / Condor
Compute Linux / Condor
Data Export 2-Person Sign-off
Data Import Scan / Move
Remote Desktop Windows
Project B
File Services CIFS / Git
Project
Remote Desktop Windows
Data Import Scan / Move
Data Export 2-Person Sign-off
REED VPN
Purdue Users
Researcher for A & B
Researcher for B
Remote Sponsors
Sponsor A
Sponsor B

CONNECTING TO REED

- Requires access to Purdue VPN
- Two-factor authentication
- FIPS 140-2 validated in-transit encryption
- Limited to (1) session
- Allows system anti-virus and other security applications to call home
- Access can be controlled geographically

ACCESSING REED

- Separate Active Directory from Campus
  - Non-US persons control, smaller user base
- Microsoft Terminal Server
  - Remote Desktop interface
  - Port-forwarding disabled
  - Running USGCB-aligned group policies
- Identified problem: Can’t print hardcopy

RUNNING ON REED

- Batch computing is an HTCondor pool, with Linux workers and Windows submitters
- Basic IP security within the pool

DATA-AT-REST & KEY MANAGEMENT

- AWS EBS containers encrypt data using FIPS 140-2 validated encryption at AES-256.
- Key management is maintained by Amazon
- Key override is AWS requires quorum and access would show in CloudTrails

AUDIT & MONITORING

- Logs routed to Audit environment
- Parsed by Splunk
- Alerts reviewed by SOC Analyst
- Can only be accessed by Audit personnel

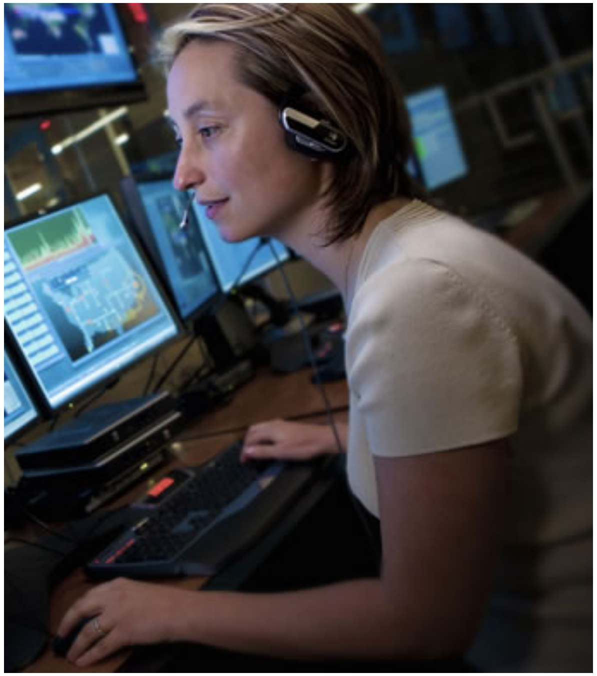

DATA INGRESS

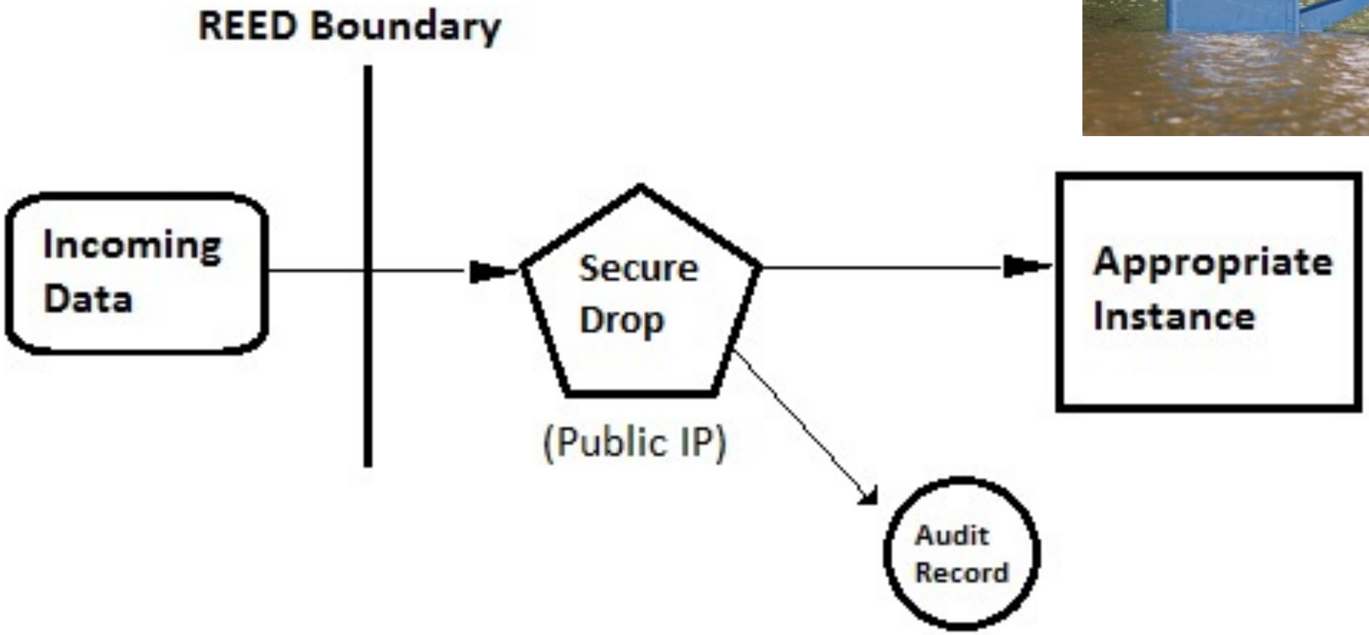

- Credentials are issued by Purdue
- Large transfers are coordinated through Purdue IT
- Data is encrypted in-transit

DATA EGRESS

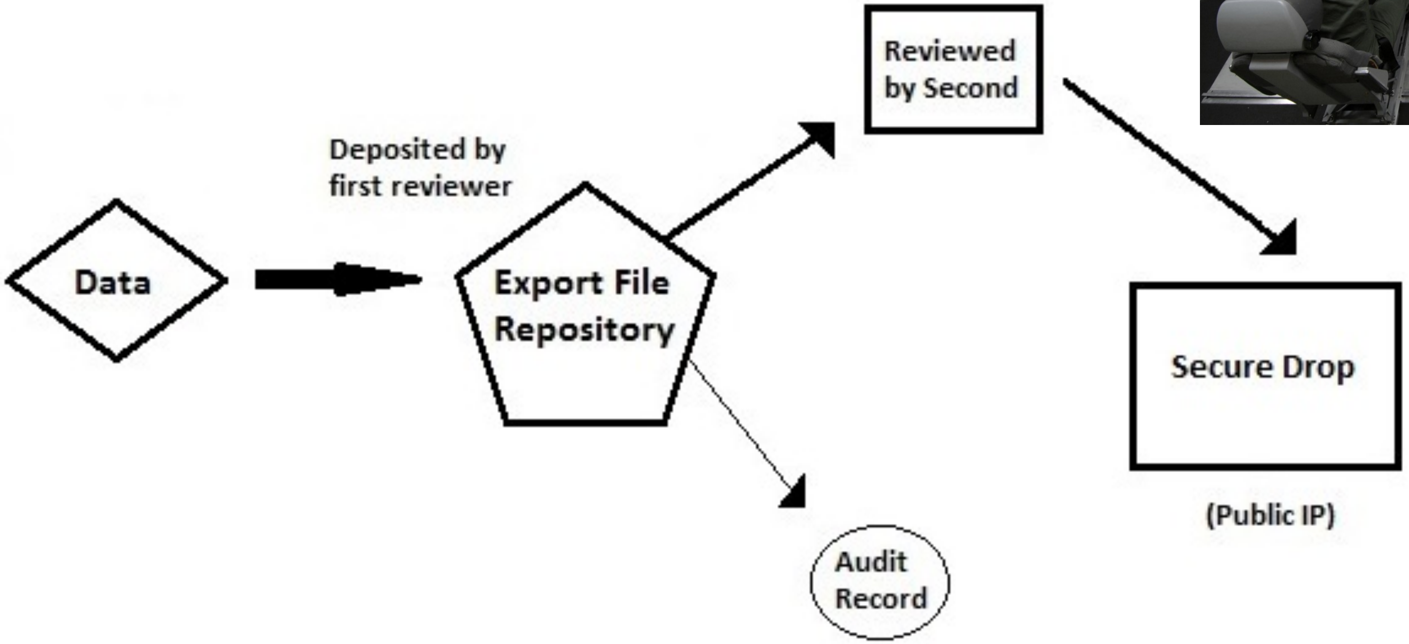

- Two-person review of data (controlled by scripts)
- Drop credentials issued by Purdue IT
- Data is encrypted in-transit

DATA ARCHIVAL

- Data storage in the cloud has cost
- Far more cost-effective to store on-premises
- Have to maintain FIPS 140-2 encryption
- Plan to use Vormetric – store on premises until recalled for new project

FUTURE WORK

BUSINESS CHALLENGES

- Expensive just to start
  - Lots of FTE cost to audit and support
- Contracts requiring REED are not large
- Sticker shock to researcher absorbing cost of compute node instances

AUTOMATION AND ONBOARDING

- Many processes still Sysadmin-driven
- Automate enrollment/disenrollment
- Automate resource assignment
  - Ala-carte style menu of options
  - Scaling based on need
- Continue to refine Log parsing & alerts

THANK YOU

Questions/Comments:

Preston Smith

psmith@purdue.edu

REED Overview credit: Jason Stein, CISSP

See our Educause Whitepaper on REED:
Leveraging Cloud Services for NIST SP 800-171, 2016

THANK YOU

Questions/Comments:

Jason Stein, CISSP
stein21@purdue.edu

See our Educause Whitepaper on REED:
Leveraging Cloud Services for NIST SP 800-171, 2016