WEB APPLICATION MONITORING AND ANALYTICS WITH SPLUNK

INFORMATION TECHNOLOGY
UNIVERSITY of WASHINGTON

AGENDA

- About Us
- What is Splunk?
- Splunk at the University of Washington
- Supporting an existing service
- Providing data to UX with client-side instrumentation
- Get Splunk for your department

ACADEMIC AND COLLABORATIVE APPLICATIONS

- A division within UW-IT focused on building student facing Web applications
- Must develop new applications while maintaining legacy applications with limited resources
- Facts and figures
  - Small team of 6 engineers
  - Maintain ~15 applications
  - Support over 140,000 users across 3 campuses
  - Support 9 groups on campus running their own Splunk instances via our license master

WHAT WE MAINTAIN

canvas

catalyst

PANOPTO

SPACESCOUT™

UNIVERSITY of WASHINGTON
IT CONNECT
Information technology tools and resources at the UW

NOTIFY.UW

MY UW

MY BACKGROUND AND ROLE

- Stephen De Vight
  - With the UW since 2006
  - Current Role: Senior Computer Specialist, 2011
  - Mission: To support teaching and learning on campus through the development of interactive Web and mobile applications

WHAT IS SPLUNK?

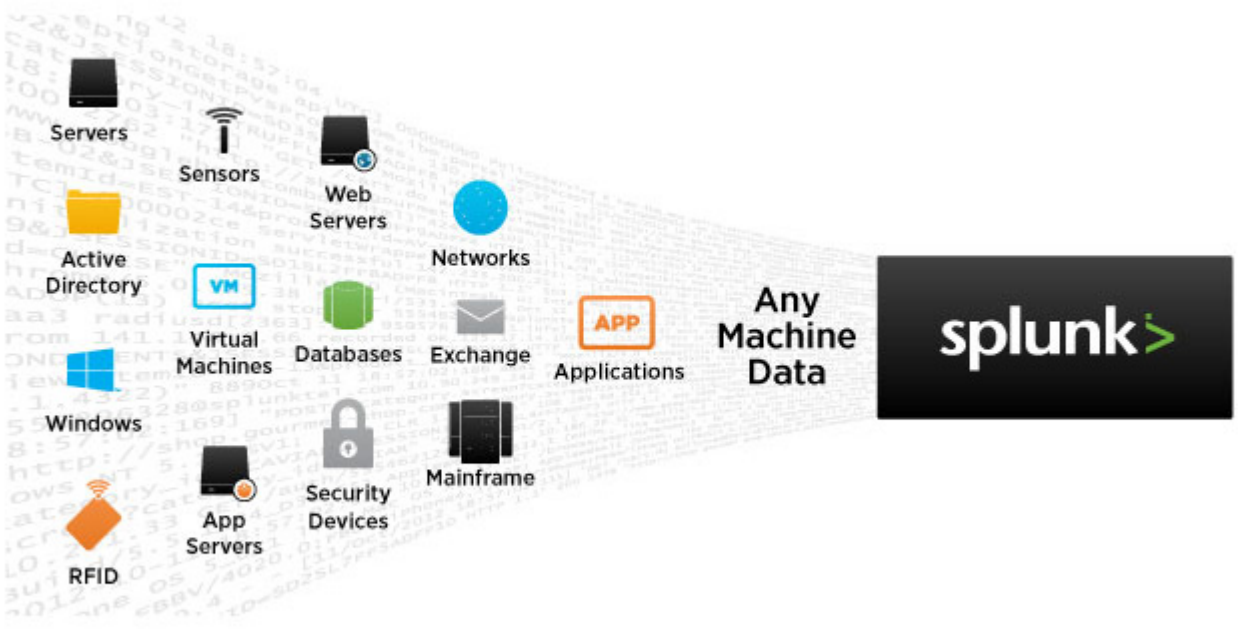

SPLUNK ENTERPRISE AT UW – 2012

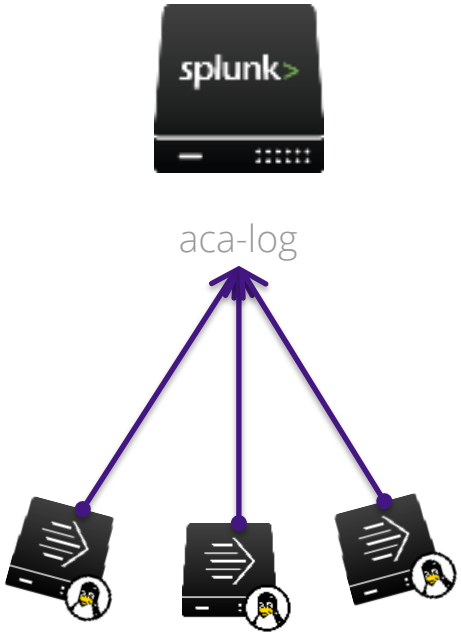

SPLUNK ENTERPRISE AT UW – 2014

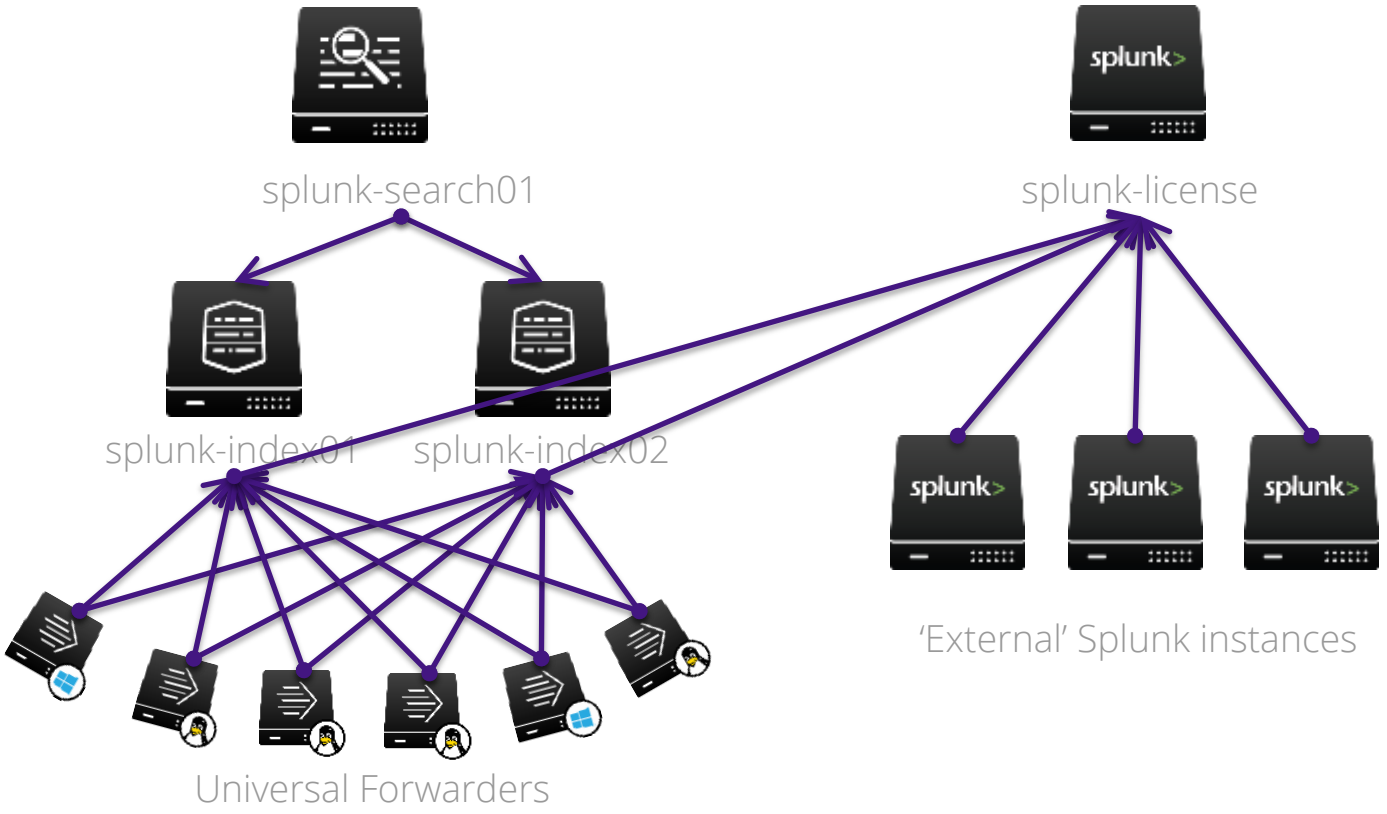

SUPPORTING AN EXISTING SERVICE

catalyst

- Homegrown suite of academic applications
- Currently consists of 8 distinct tools
- Released in 1999

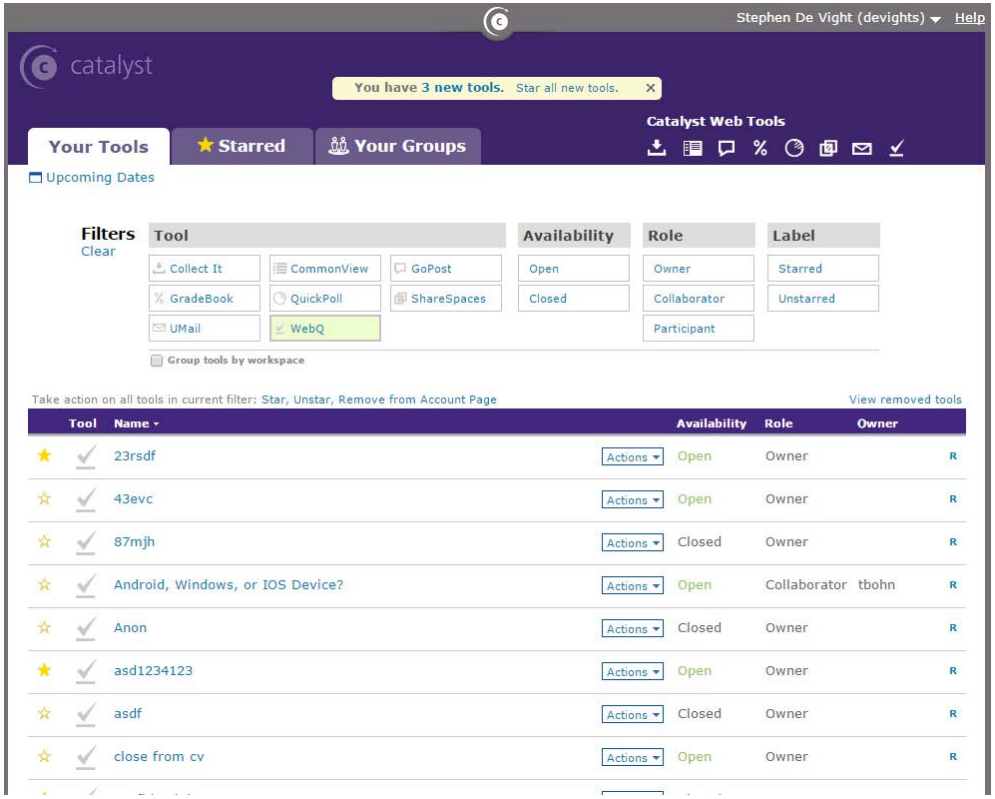

SUPPORTING AN EXISTING SERVICE

OUR NEEDS

> Situation: Legacy database logging system reached end of life, was not scaling well, and was too costly to directly replace

> Struggling with: Finding a solution that is both easy to build and maintain as well as being able to scale to our needs

> Wanted: An easy to use, UI-driven, application to search our log data

> Enter Splunk: Splunk Enterprise allowed us to build a custom searching app as well as a dashboard for monitoring service status

SUPPORTING AN EXISTING SERVICE

CATALYST LOG SEARCH

> Splunk application with advanced XML view

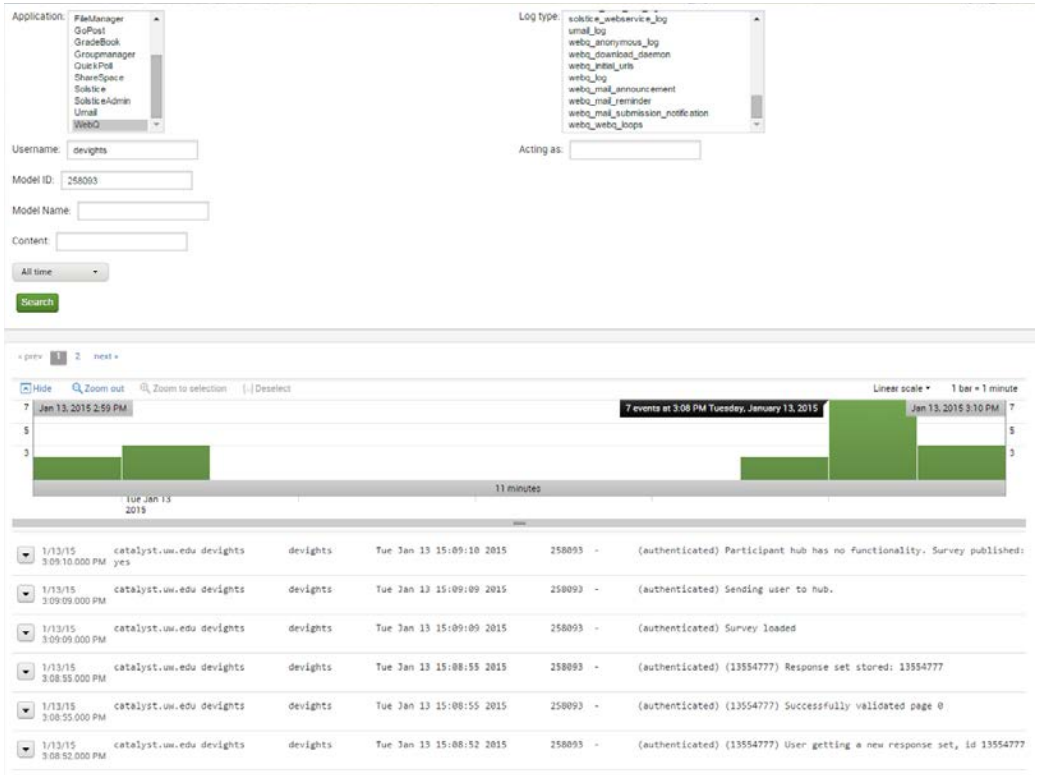

SUPPORTING AN EXISTING SERVICE

CATALYST LOG SEARCH

- Splunk application with advanced XML view
- Search form negates the need for users to learn Splunk search language or understand our log formatting and structure

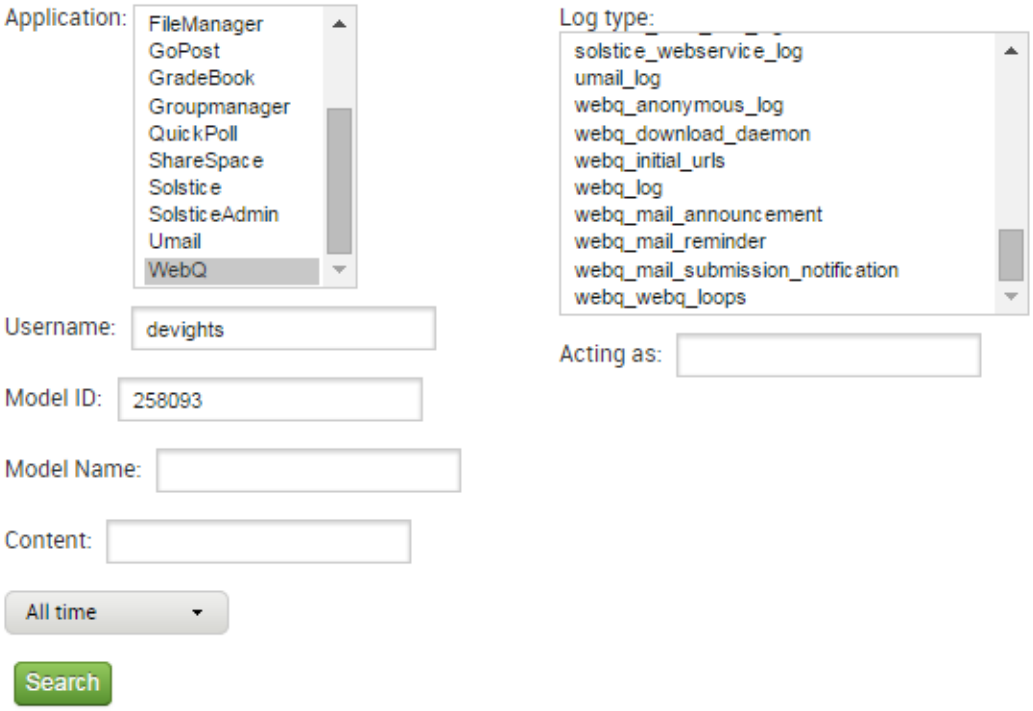

SUPPORTING AN EXISTING SERVICE

CATALYST LOG SEARCH

- Splunk application with advanced XML view
- Search form negates the need for users to learn Splunk search language or understand our log formatting and structure
- Support can analyze user activity to provide insight into incident reports

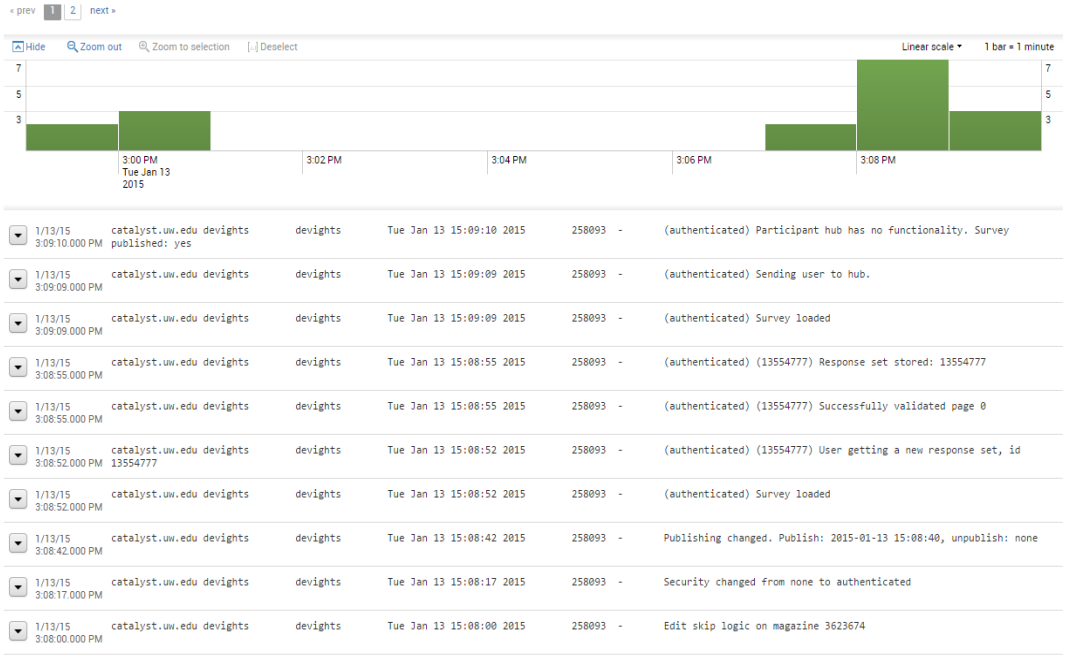

SUPPORTING AN EXISTING SERVICE

CATALYST DASHBOARD

- Gauge current level of activity at a glance
- Examine last day of activity for anomalous usage
- Targets slowest loading URLs for performance improvement

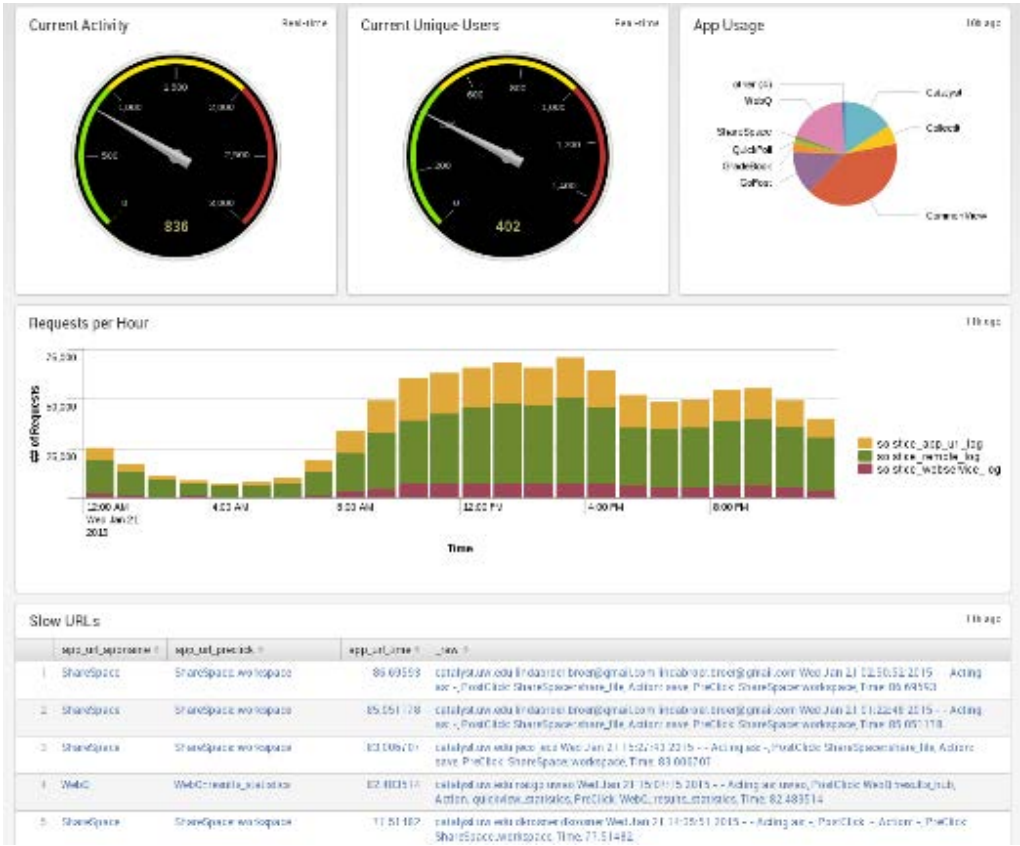

DATA DRIVEN USER EXPERIENCE

- Mobile Web version of our student portal
- Focused on providing timely, actionable information to our students
- Based on a student's situation and the time of the quarter we dynamically display, hide, move, and reorder content

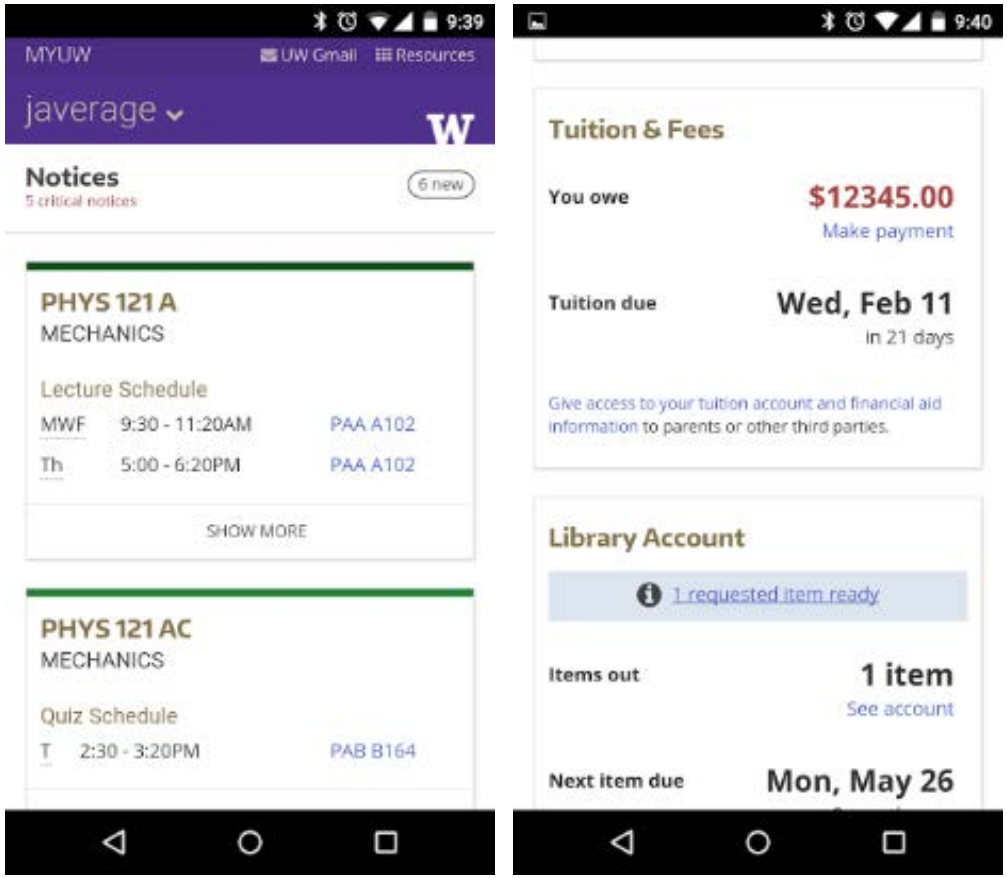

DATA DRIVEN USER EXPERIENCE

OUR NEEDS

> Situation: UX needs a way to validate their assumptions around what content is relevant to a student at various points in the quarter

> Struggling with: Correlating user activity with institutional data (e.g. class standing, campus, etc.)

> Wanted: A self-driven means for UX and business analysts to analyze log data

> Enter Splunk: Splunk, along with our client-side logging solution, allows us to correlate user activity with certain institutional attributes we log

DATA DRIVEN USER EXPERIENCE

CLIENT-SIDE LOGGING

> Google Analytics did not get us everything we needed
> Using log4javascript to collate events and POST to a REST interface
> Events are bundled to reduce network overhead
> Events are written to file by REST server

http://www.log4javascript.org/

Name Path	Method	Status Text	Initiator	Size ▲ Content	Time Latency	Timeline 10.00 s
log /logging	POST	200 OK	log4javascript.js:2211 Script	176 B 0 B	16 ms 14 ms	
log /logging	POST	200 OK	log4javascript.js:2211 Script	176 B 0 B	17 ms 15 ms	
log /logging	POST	200 OK	log4javascript.js:2211 Script	176 B 0 B	19 ms 17 ms	

DATA DRIVEN USER EXPERIENCE

WORKING WITH CLIENT LOGS

- Link Log
  - Link location
  - Target URL
  - Action (view, click)
- Card Log
  - Card location URL
  - Card name
  - Card position
  - Action (load, view, expand, collapse)

```
INFO 21 22:25:31
{
   "level": "INFO",
   "url": "https://my.uw.edu/mobile/landing/",
   "timestamp": 1421907930962,
   "logger": "link",
   "session_key": "xc63940325jlo3dsdfcgtt3126b",
   "message": {
       "href": "http: //gmail.uw.edu/",
       "action": "click"
   }
}
[link]
```

DATA DRIVEN USER EXPERIENCE

```
index=myuw_production
sourcetype=myuw_link_log
action=click
|stats count by target_url
```

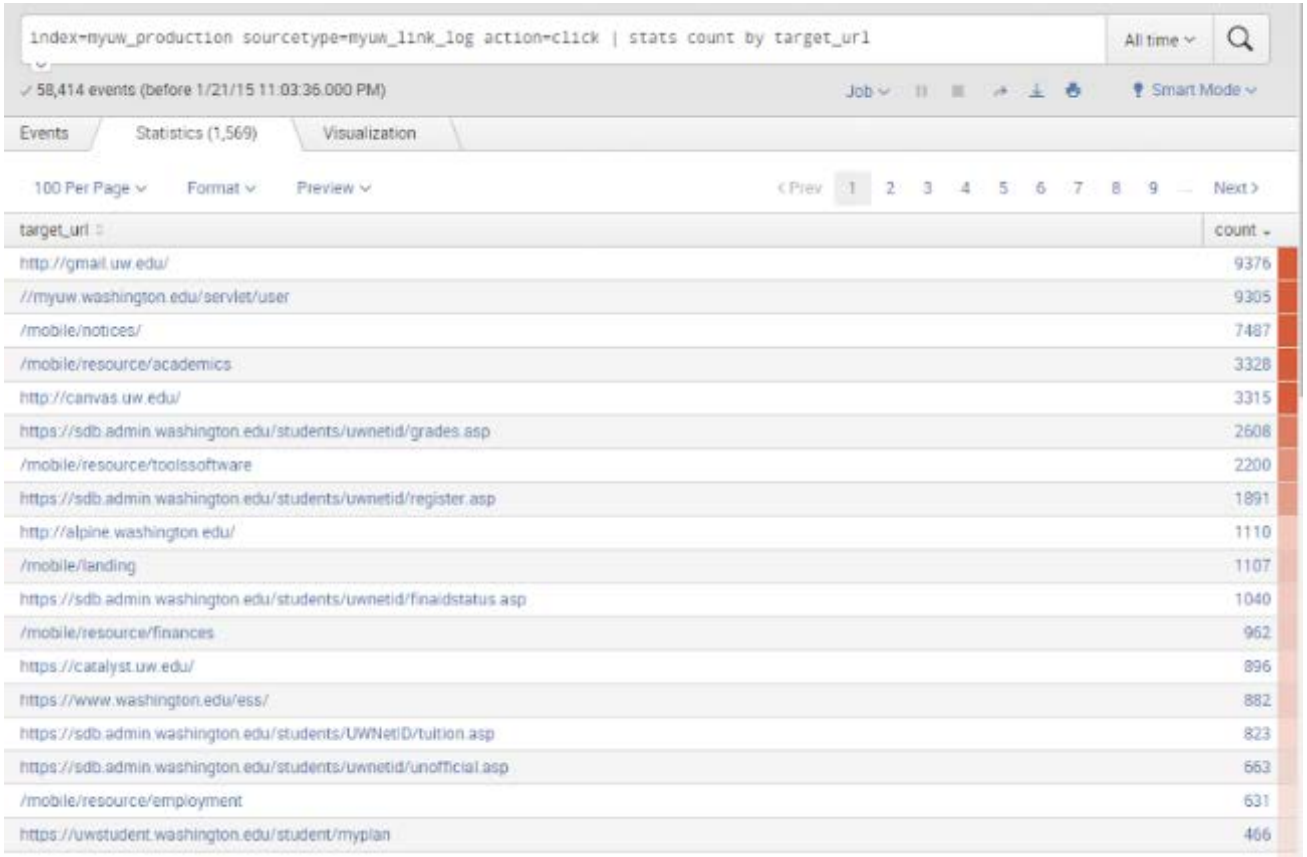

DATA DRIVEN USER EXPERIENCE

SERVER-SIDE SESSION LOG

- > Session Log
  - > Graduate or undergraduate
  - > Class standing
  - > Campus

```
INFO 21 22:21:20
{
   "is_grad": false,
   "netid": "javerage",
   "is_ugrad": true,
   "class_level": "FRESHMAN",
   "session_key": "xc63940325jlo3dsdfcgtt3126b",
   "campus": "seattle"
}
[session]
```

DATA DRIVEN USER EXPERIENCE

EVENTTYPES AND TRANSACTIONS

> Build an eventtype that contains both link and session logs

```
index=myuw_production
(sourcetype=myuw_link_log
OR sourcetype=myuw_session_log)
```

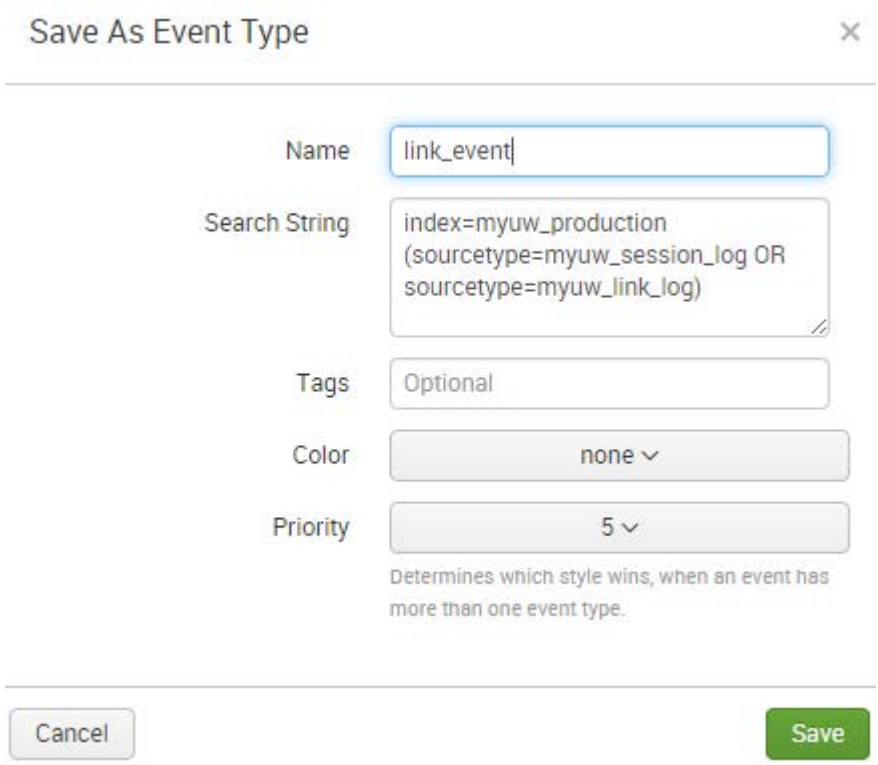

DATA DRIVEN USER EXPERIENCE

SESSION ACTIVITY WITH TRANSACTIONS

- Create a transaction based on session_key
- Find transactions that contain a link click to '*dars.asp'
- Get count of other URL targets clicked within that transaction

```
index=myuw_production
    eventtype=link_event
|transaction fields=session_key
    maxspan=8h
|search target_url=*dars.asp
    AND action=click
|stats count by target_url
```

target_url	count
https://sdb.admin.washington.edu/students/uwnetid/dars.asp	5022
https://uwstudent.washington.edu/student/myplan	4996
/mobile/landing	4861
https://sdb.admin.washington.edu/students/uwnetid/register.asp	4761
https://sdb.admin.washington.edu/timeschd/uwnetid/findschd.asp	4158
http://www.washington.edu/students/reg/calendar.html	4075
/mobile/notices/	3844
/mobile/resource/academics	3839
/mobile/resource/finances	3829
http://www.washington.edu/students/timeschd/	3784
http://www.washington.edu/students/crscat/	3757
http://careers.washington.edu/	3738
/mobile/resource/studentcampuslife	3489
/mobile/resource/eventsactivities	3488
http://depts.washington.edu/aspuw/clue/home/	3469

DATA DRIVEN USER EXPERIENCE

COMBINING LOGS WITH TRANSACTIONS

- Create a transaction based on session_key
- Find link events that have a click action
- Using the session log, determine how many link clicks were made by each class level

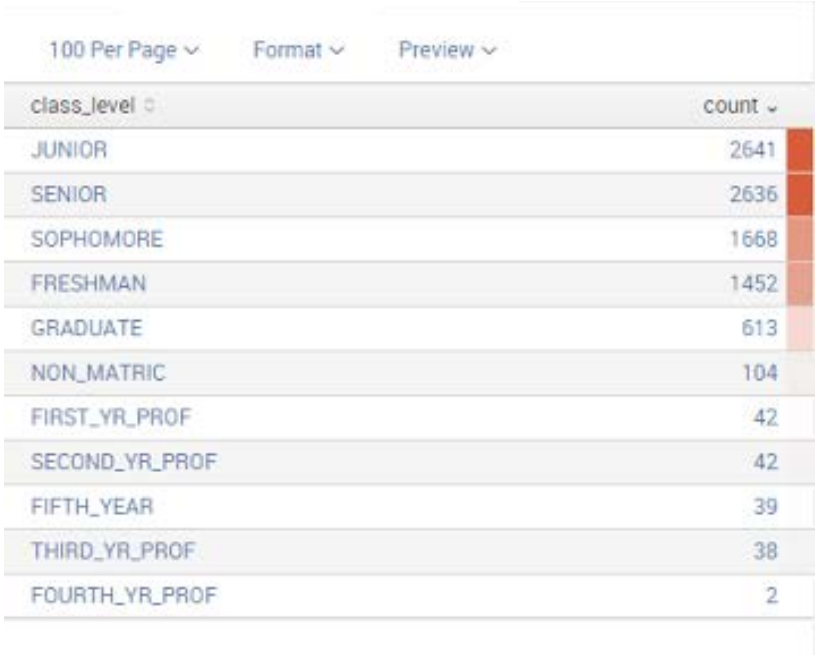


```
index=myuw_production eventtype=link_event
|transaction fields=session_key maxspan=8h
|search action=click
|stats count by class_level
```

TOP TAKEAWAYS

- Building a search form makes Splunk simple to use
- Determine your analysis needs before creating your logging scheme
- Client side logging can provide valuable insight into user behavior
- Transactions make combining logs easy

SPLUNK FOR YOUR DEPARTMENT

> Splunk is sold in terms of data indexed per day
> Discounted pricing available through Internet2
> Contact tomlewis@uw.edu for details

Tier (GB)	Fees for Software and Support (total amount, payable over three years in annual installments)	Example Annual Payment	$/GB/YR
20	$44,577	$14,859/YR	743
50	$70,581	$23,527/YR	471
100	$111,447	$37,149/YR	371
200	$208,032	$69,344/YR	347
500	$482,937	$160,979/YR	322
1,000	$891,573	$297,192/YR	290

QUESTIONS?

INFORMATION TECHNOLOGY

UNIVERSITY *of* WASHINGTON