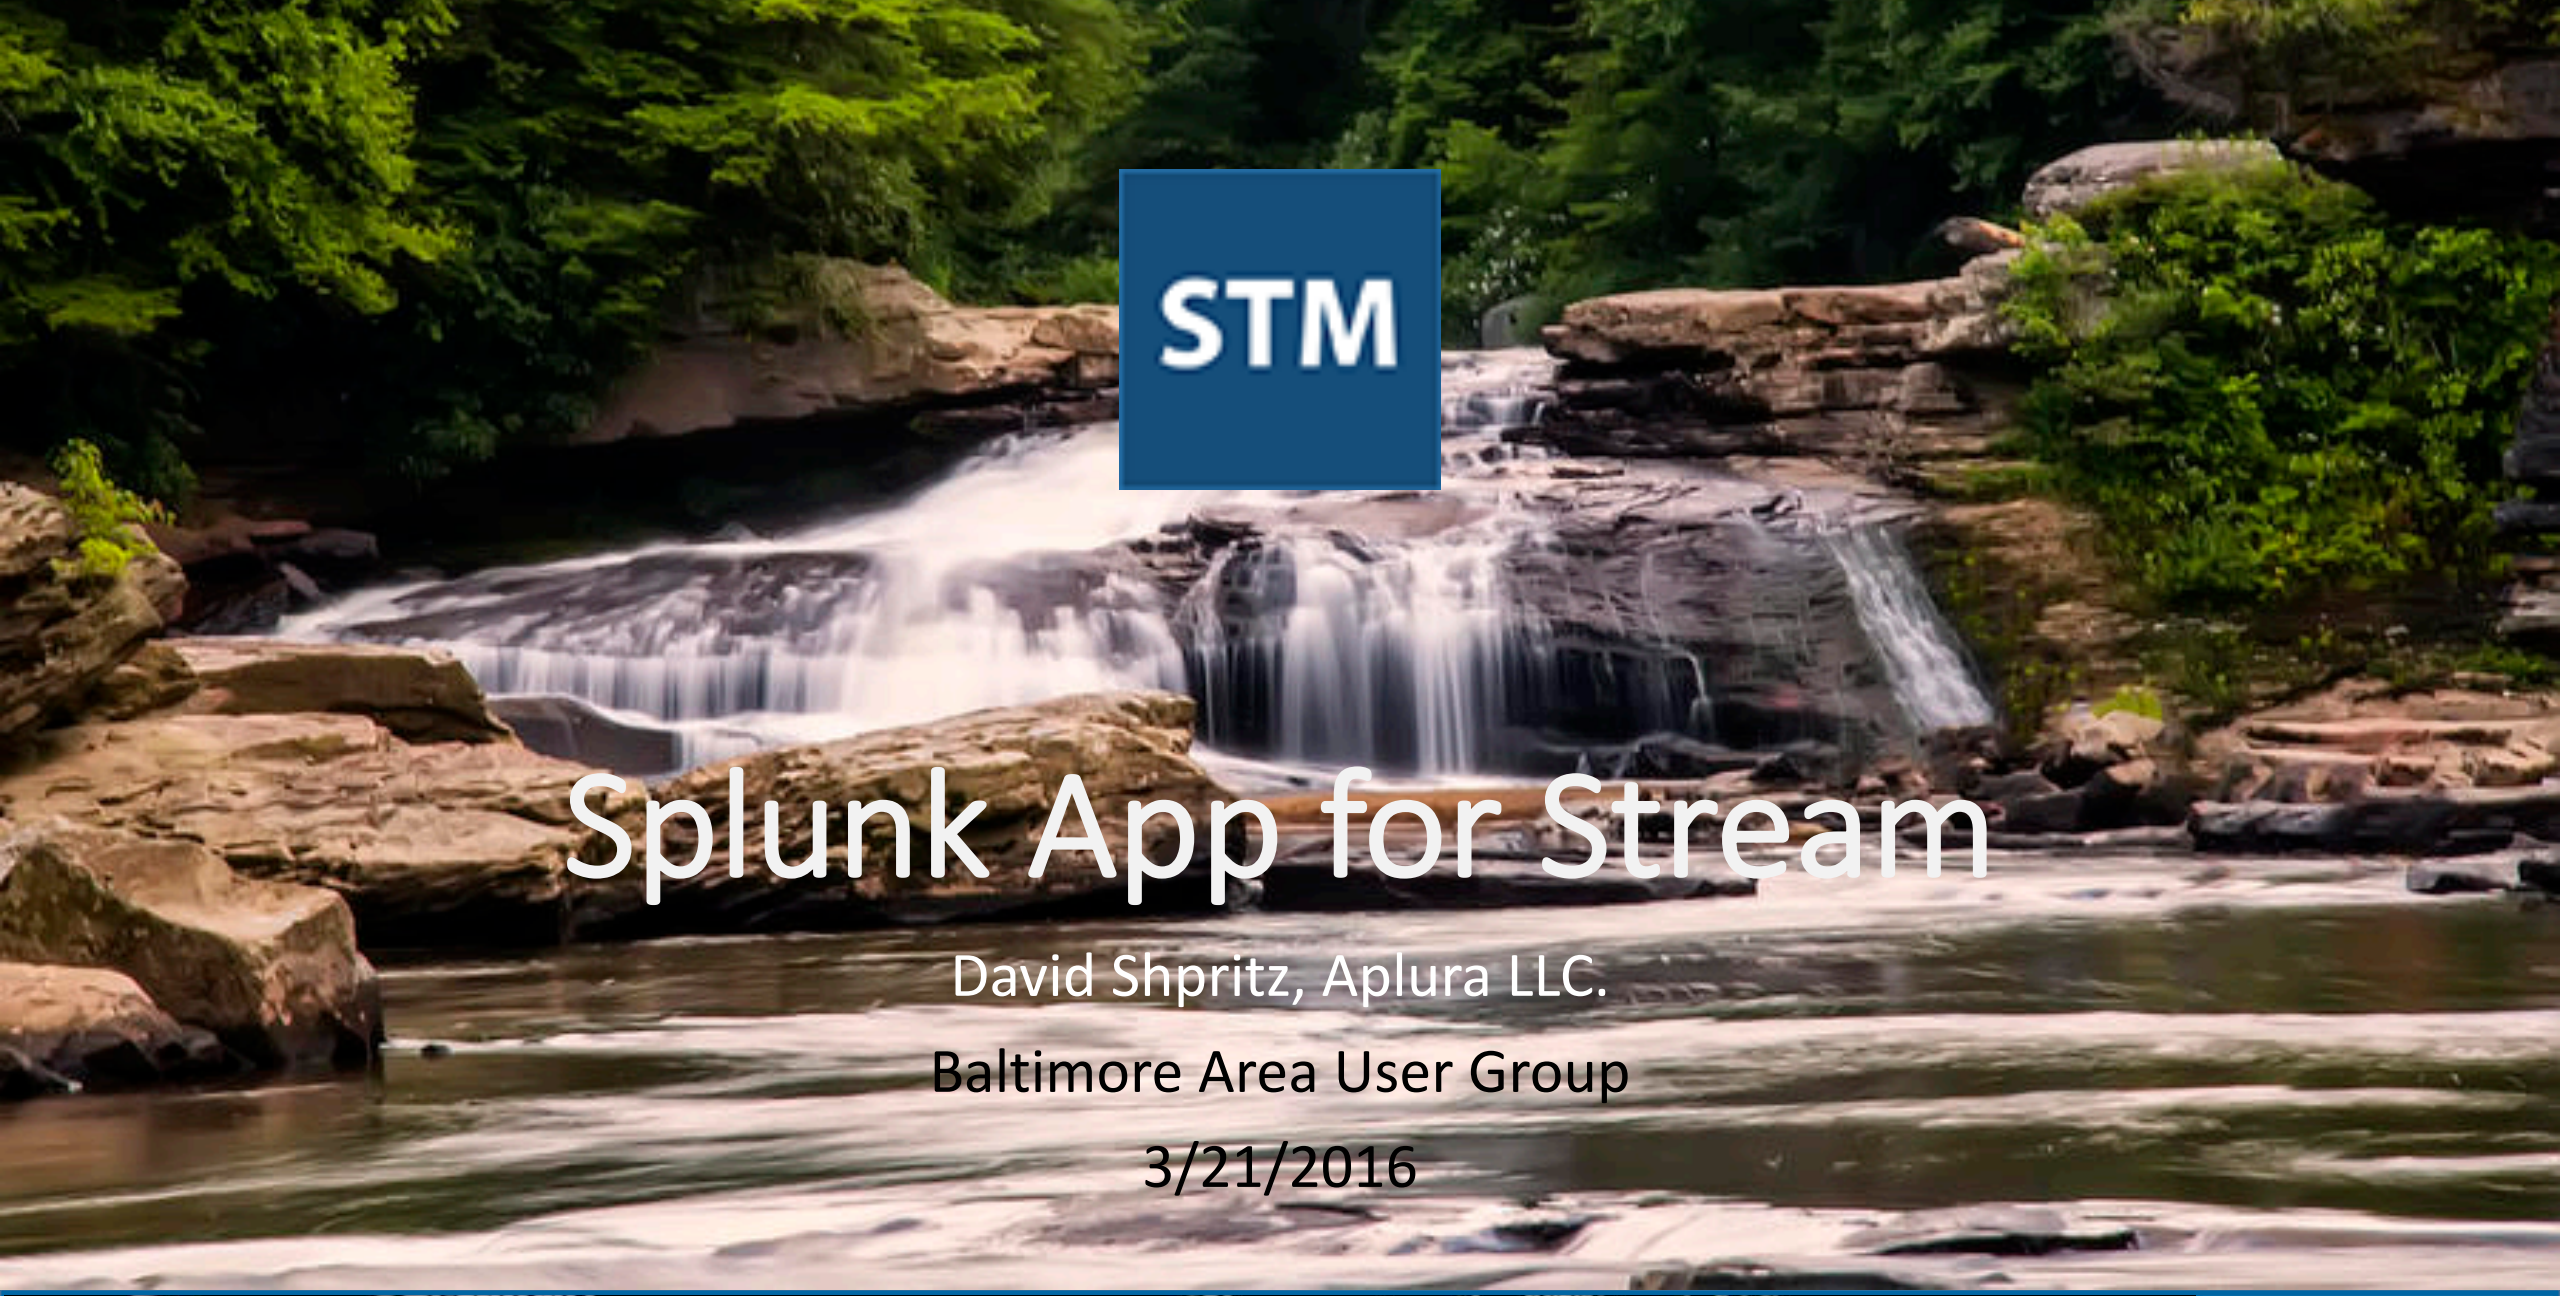

Splunk App for Stream

David Shpritz, Aplura LLC.

Baltimore Area User Group

3/21/2016

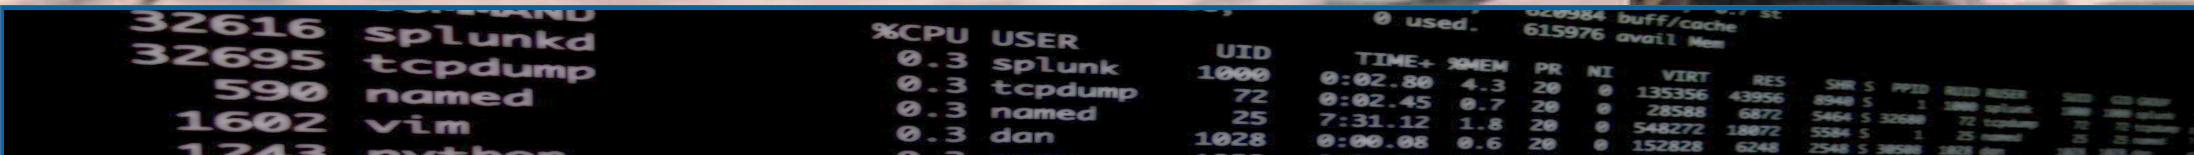

APLURA

Many Solutions, One Goal.

Agenda

- What is Splunk App for Stream?
- Why use Steam?
- Where to use Stream?
- Deploying Stream
- Questions

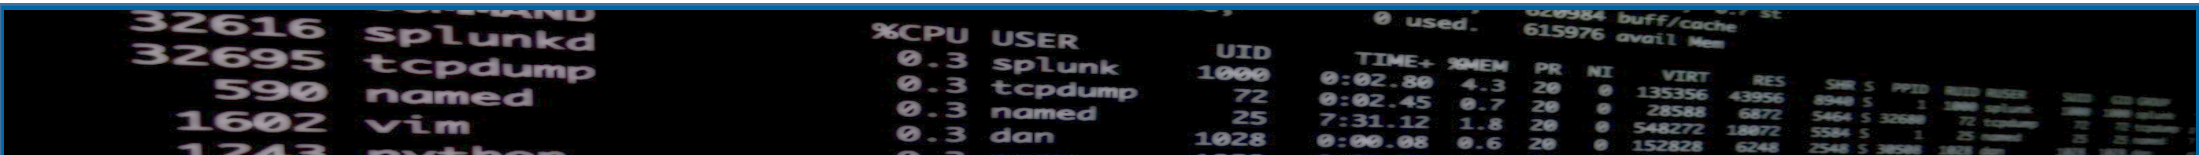

Many Solutions, One Goal.

What Is Splunk App for Stream?

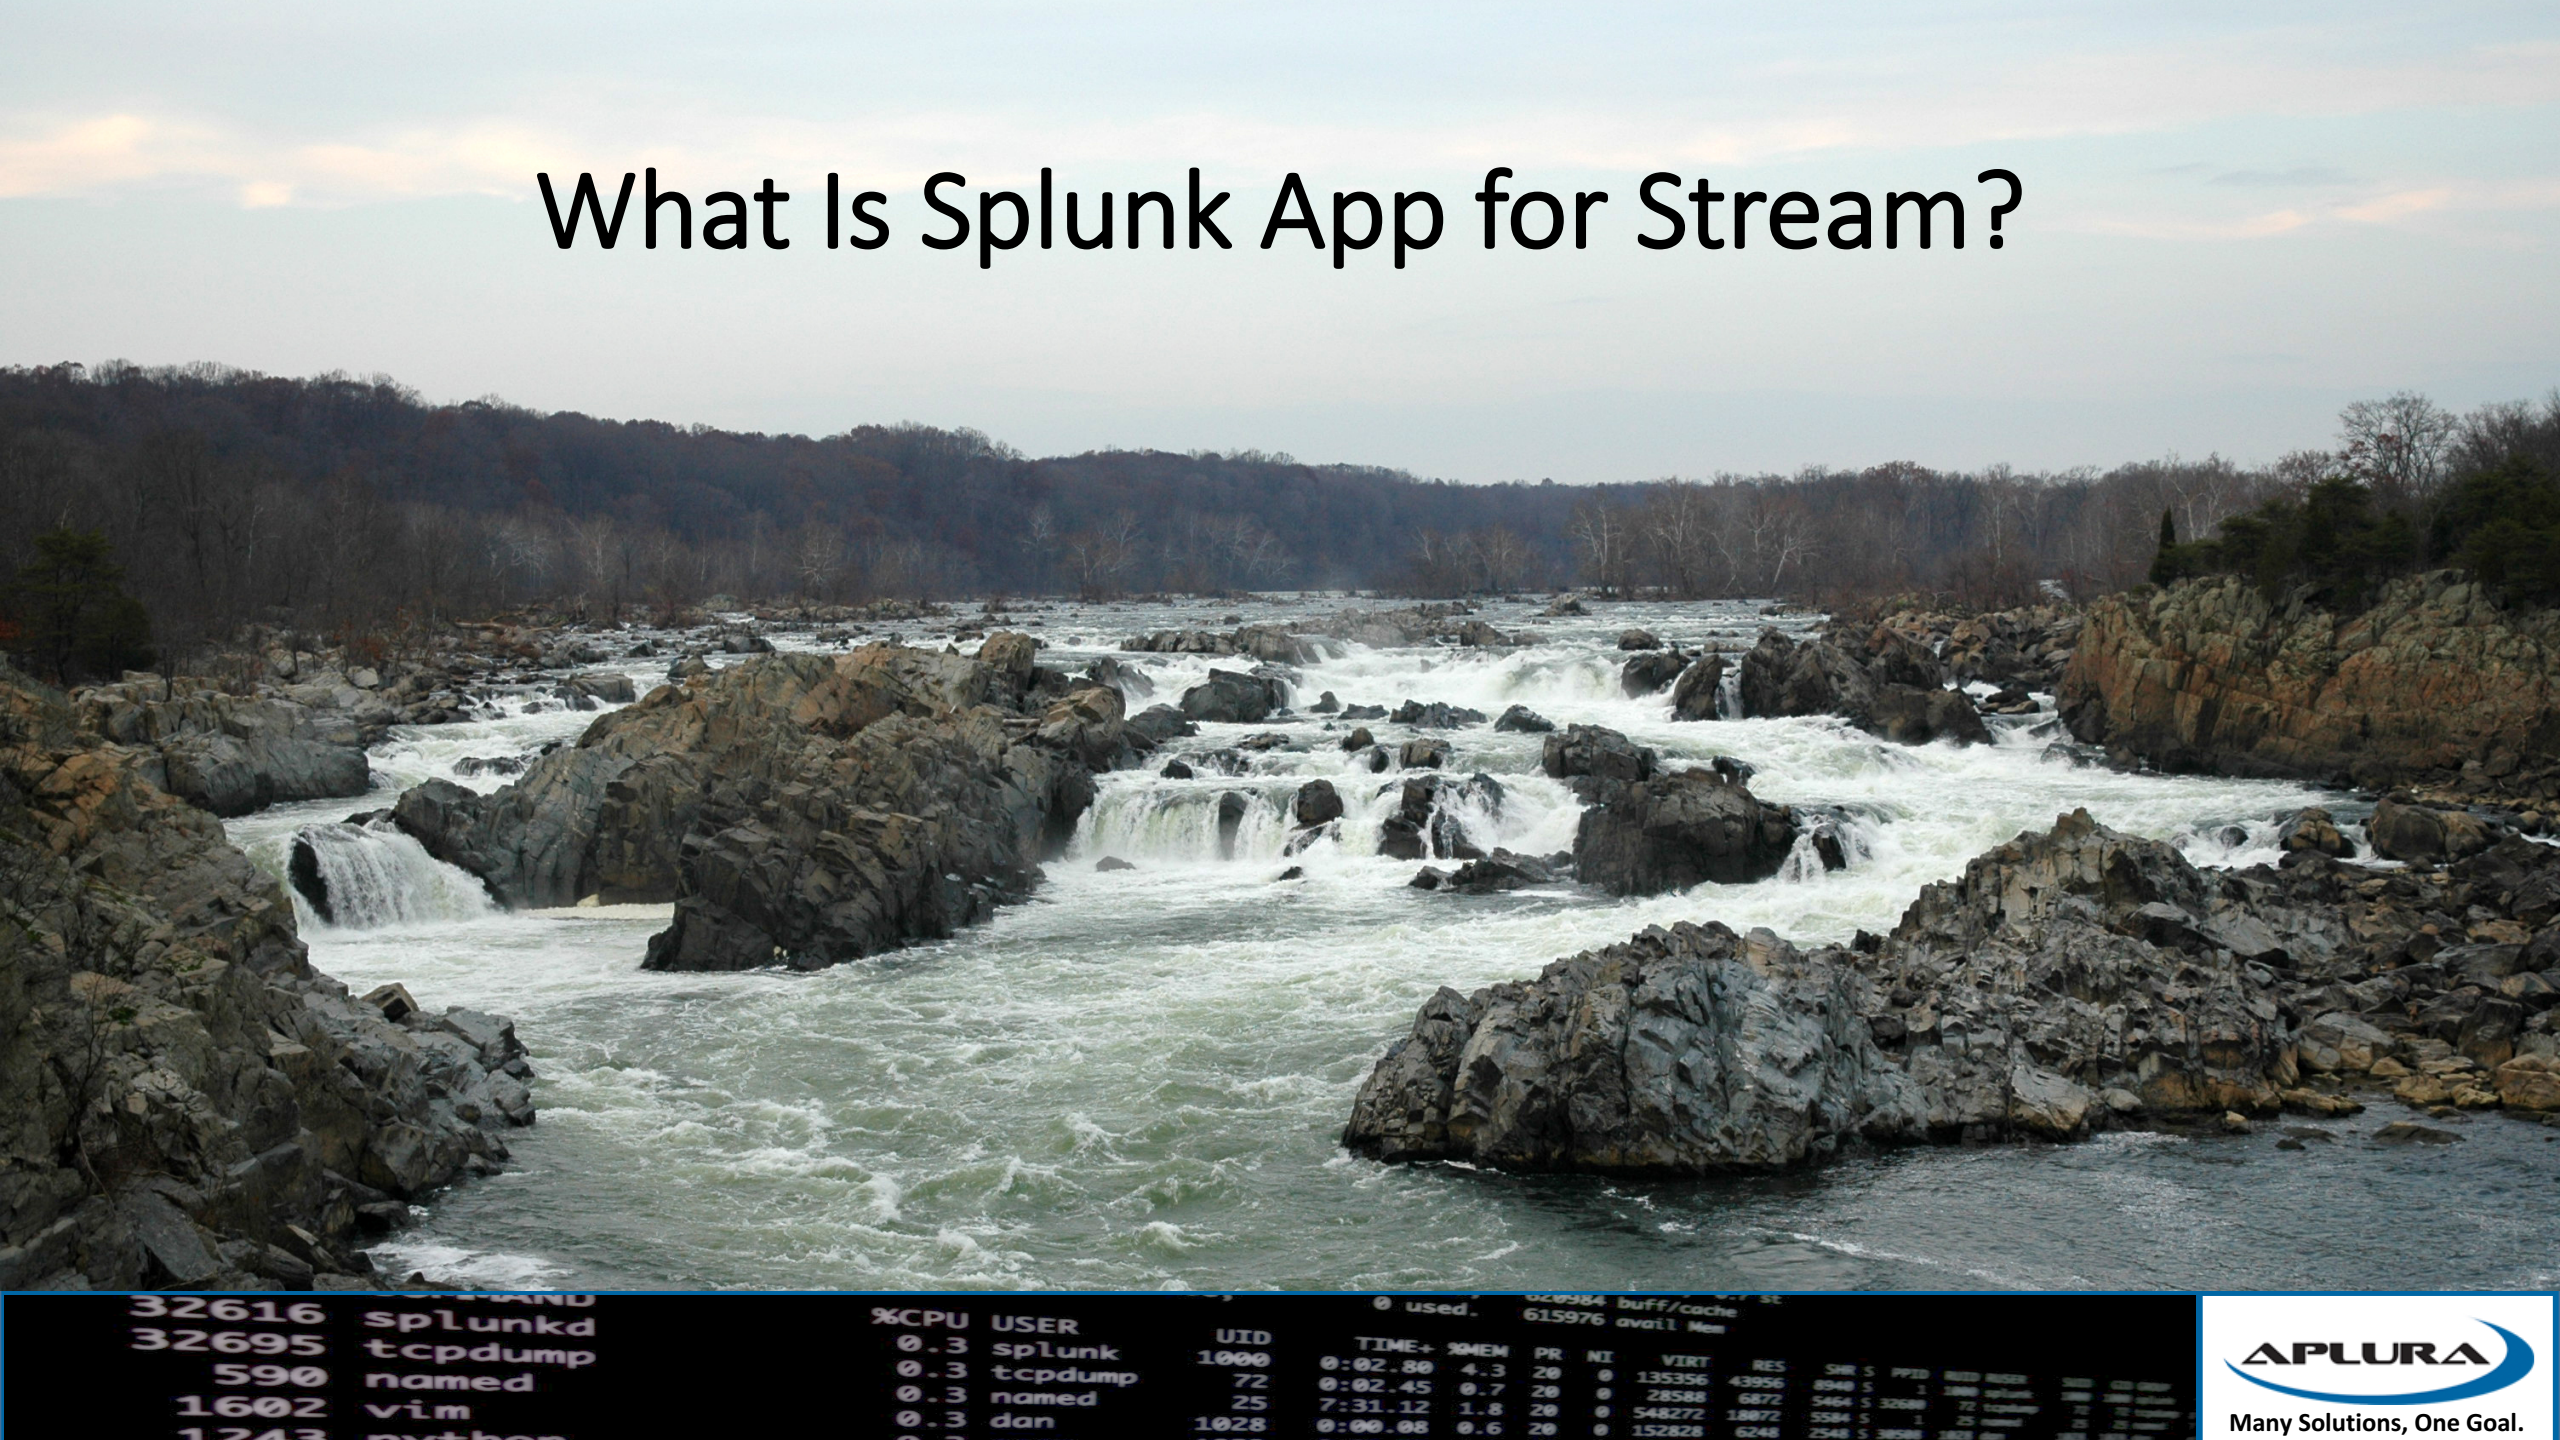

APLURA

Many Solutions, One Goal.

Some history

- Splunk acquires Cloudmeter, December 2013
- Renamed Splunk App for Stream
- Released with Splunk 6.0 (August, 2014)
- Now at version 6.4.3 (January, 2016)

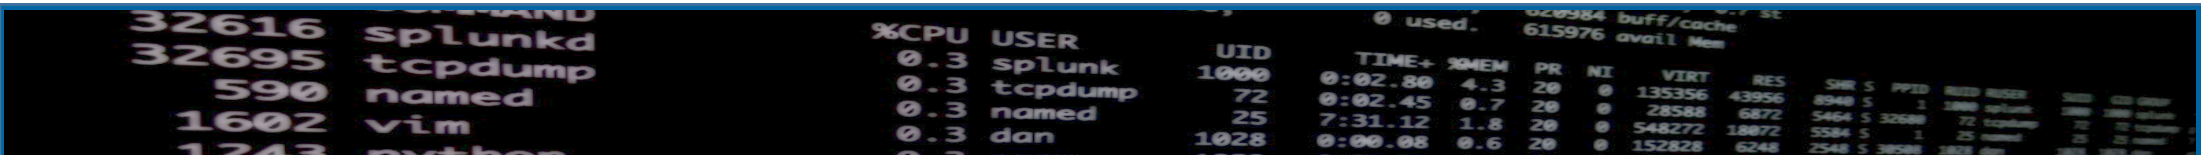

Purpose of Stream

- Rapid deployment
- Rapid configuration
- Capture wire data
- Interpret wire data
- Summarize/filter/aggregate
- Index
- Kind of like Bro, but more Splunky, and GUI

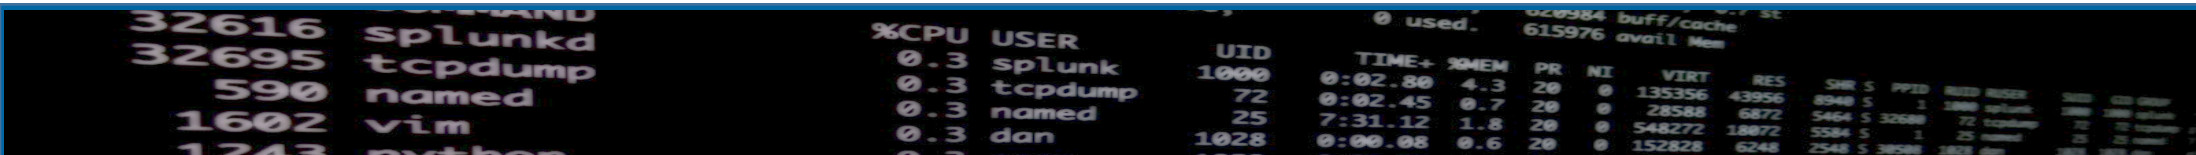

So what can we capture?

- Well, we aren't really capturing and indexing packets
- Forwarders capture packets, analyze the protocols
- What protocols (a lot):
  - TCP/UDP
  - Application protocols (HTTP, databases, email, file sharing, chat)
  - About 30 different protocols currently
  - http://docs.splunk.com/Documentation/StreamApp/latest/DeployStreamApp/Whattypeofdatadoesthisappcollect

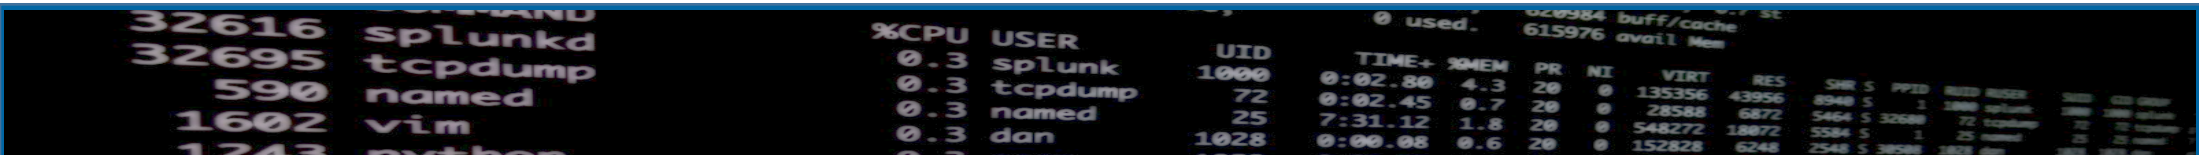

Why to use Splunk Stream

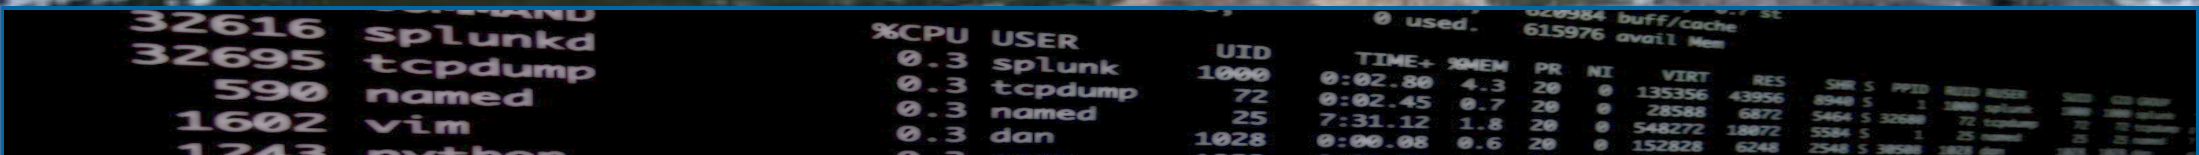

APLURA

Many Solutions, One Goal.

No logs

- No ownership
- No visibility
- No forwarders (as endpoints)
- No logging options

MISSING HOMEWORK LOG

Name:_____ Class:_____

Date	Assignment Name	Reason why the assignment is not here today	Next Steps Check the column which applies: I will have it done for tomorrow	I will come in tomorrow at 7:30	Check here when completed

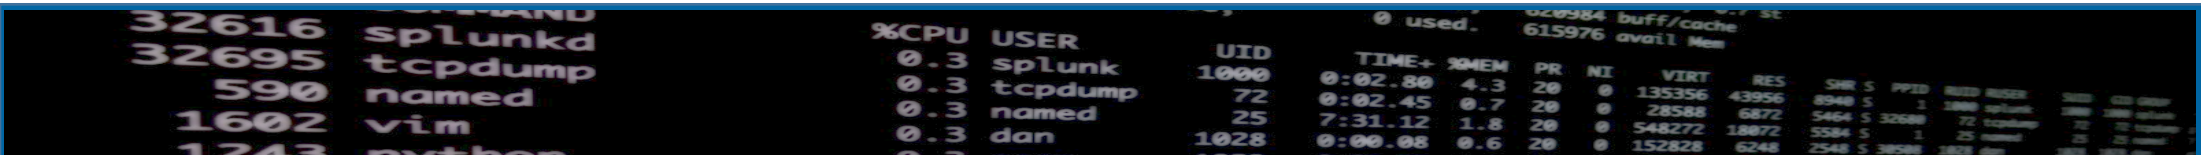

APLURA

Many Solutions, One Goal.

Poor logs

- Logging is high overhead
- Logs make no sense
- Key events are not logged

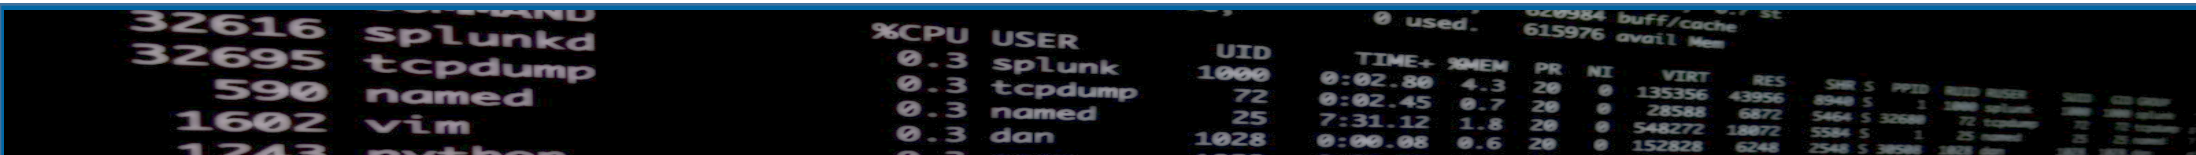

APLURA
Many Solutions, One Goal.

Cloud

- Many cloud services don’t offer logs on things
- No chokepoints

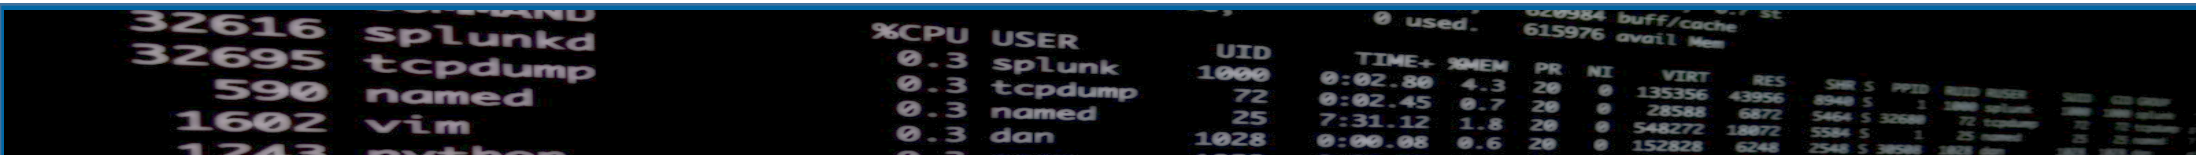

APLURA
Many Solutions, One Goal.

VS. Bro IDS

- Lower CPU usage
- Lower RAM usage
- More OS support (Linux, Windows, OSX)

But

- High traffic requires network packet brokers (Gigamon, Ixia, etc.)
- Can’t write your work interpreters
- No Snort rules

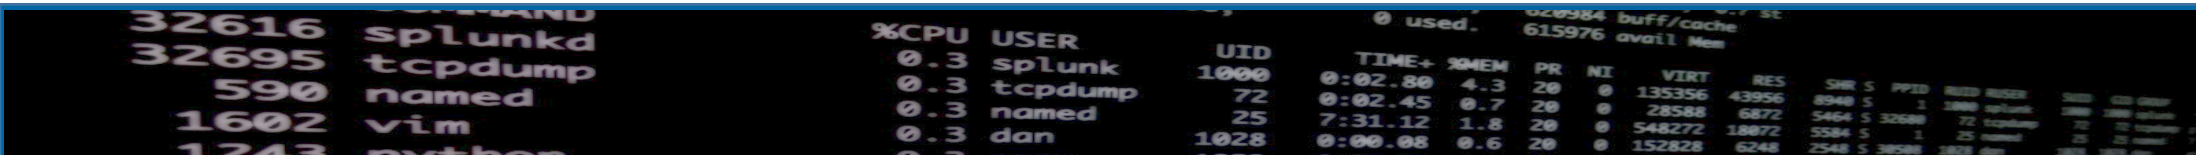

Other features

- Filtering
- Aggregation
- Ephemeral Streams (short term)
- SSL decrypt
- Centralized management
- Integration with ES
  - Start a stream after Notable event
  - Protocol analysis dashboards

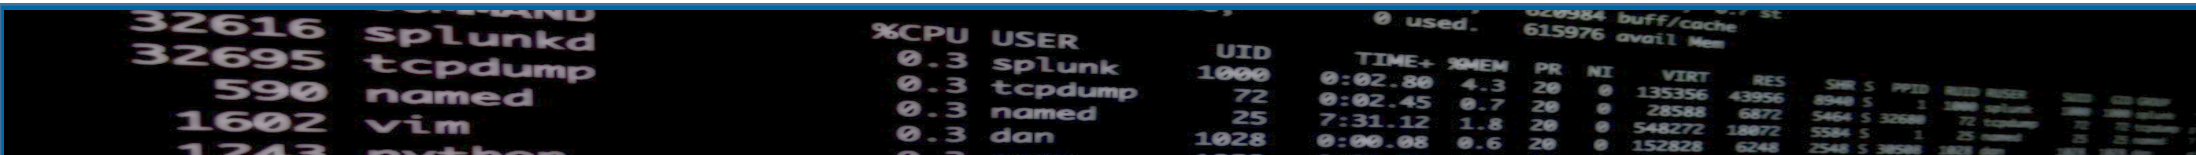

APLURA
Many Solutions, One Goal.

Data Estimation

- “What if I turn this on?”
- Tells you how much data you would be indexing

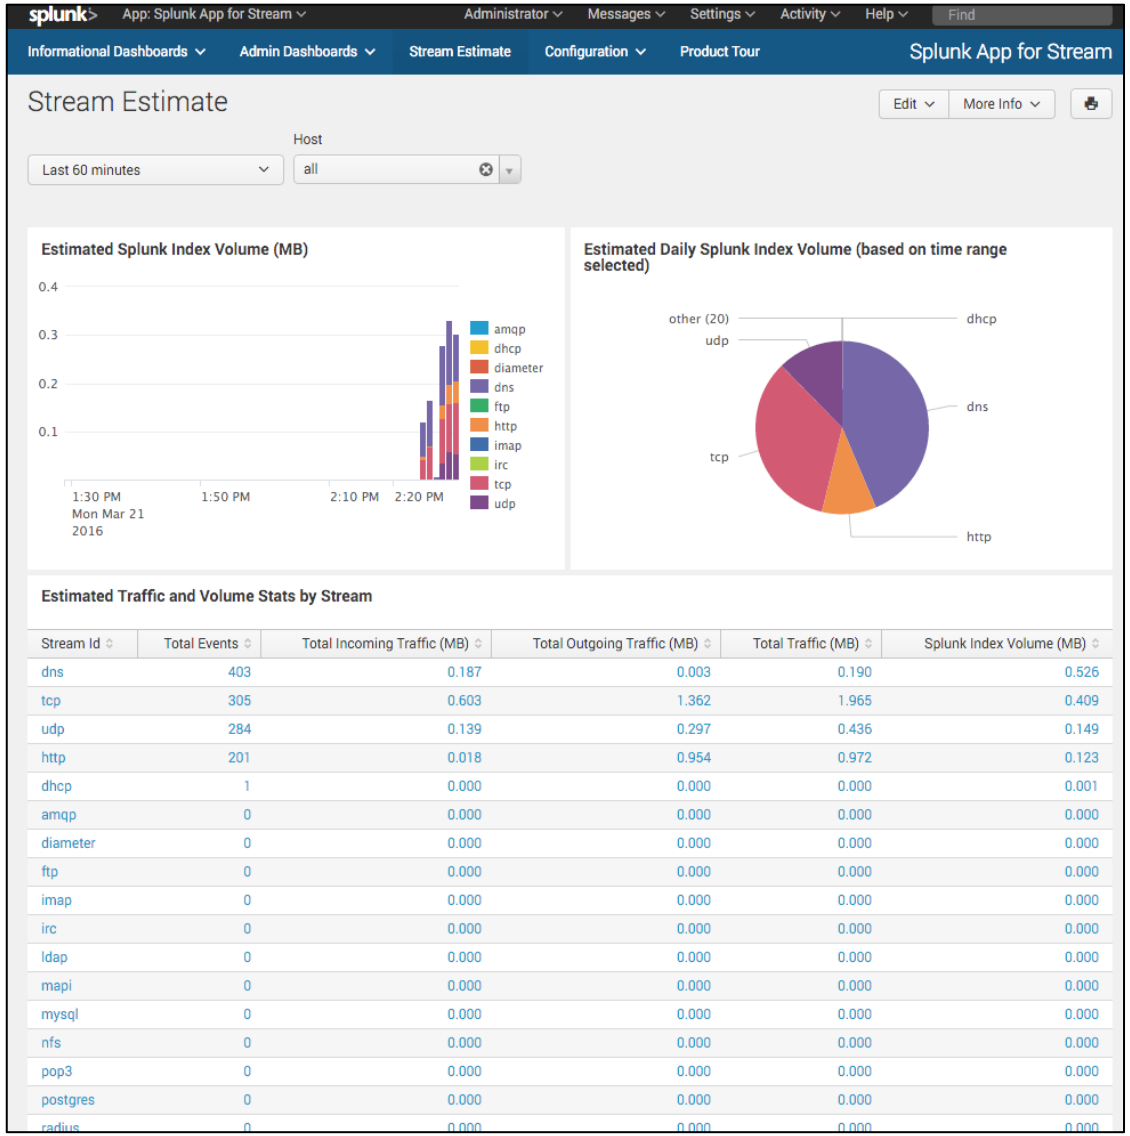

Granular control of the data

- Not just which systems, but also what data, which fields

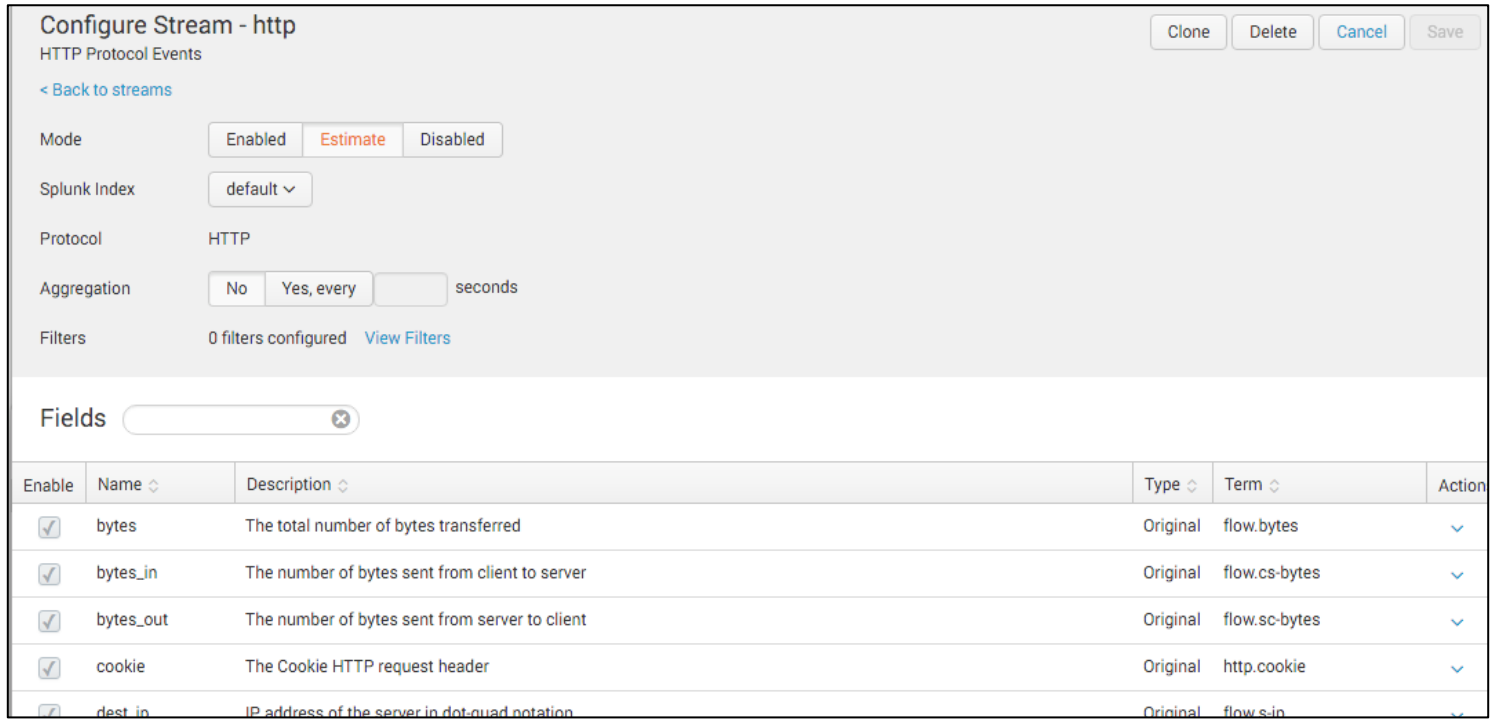

Global Filters

- Filter out noise from the enterprise
- Things like vulnerability scanners

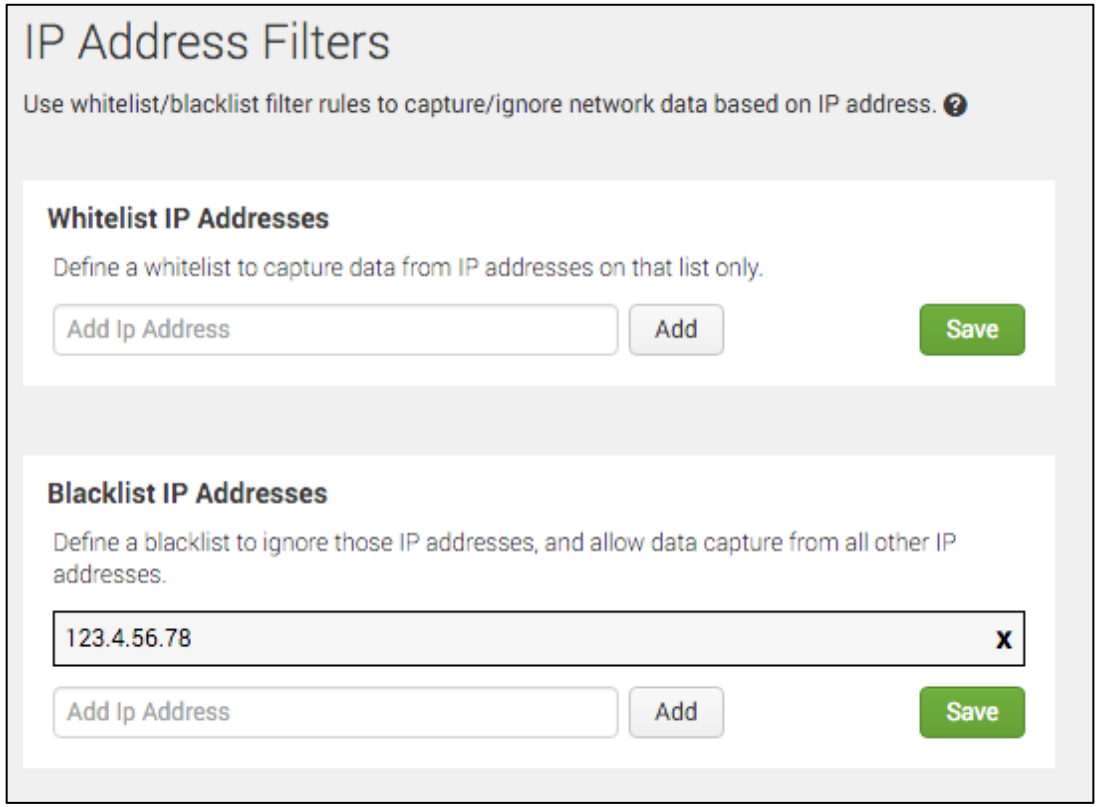

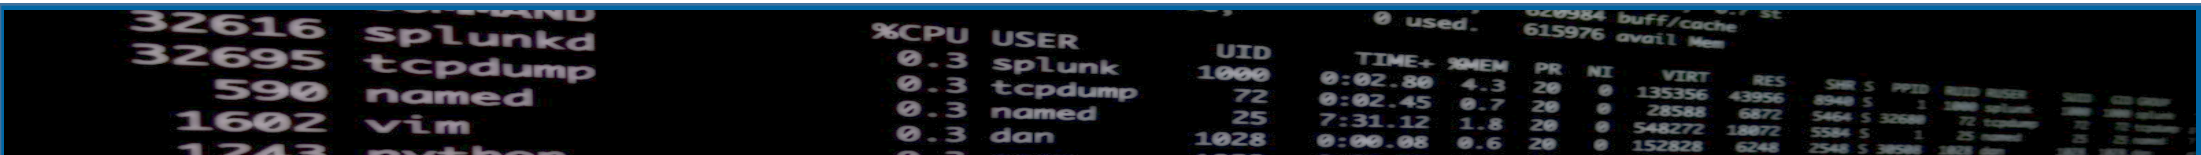

APLURA
Many Solutions, One Goal.

Distributed Forwarder Management

- Set up groups for capture
- Uses regex for groups on the “Forwarder ID”
- Forwarder ID is configurable via XML config file
- Yes, it’s another Splunk deployment/control mechanism

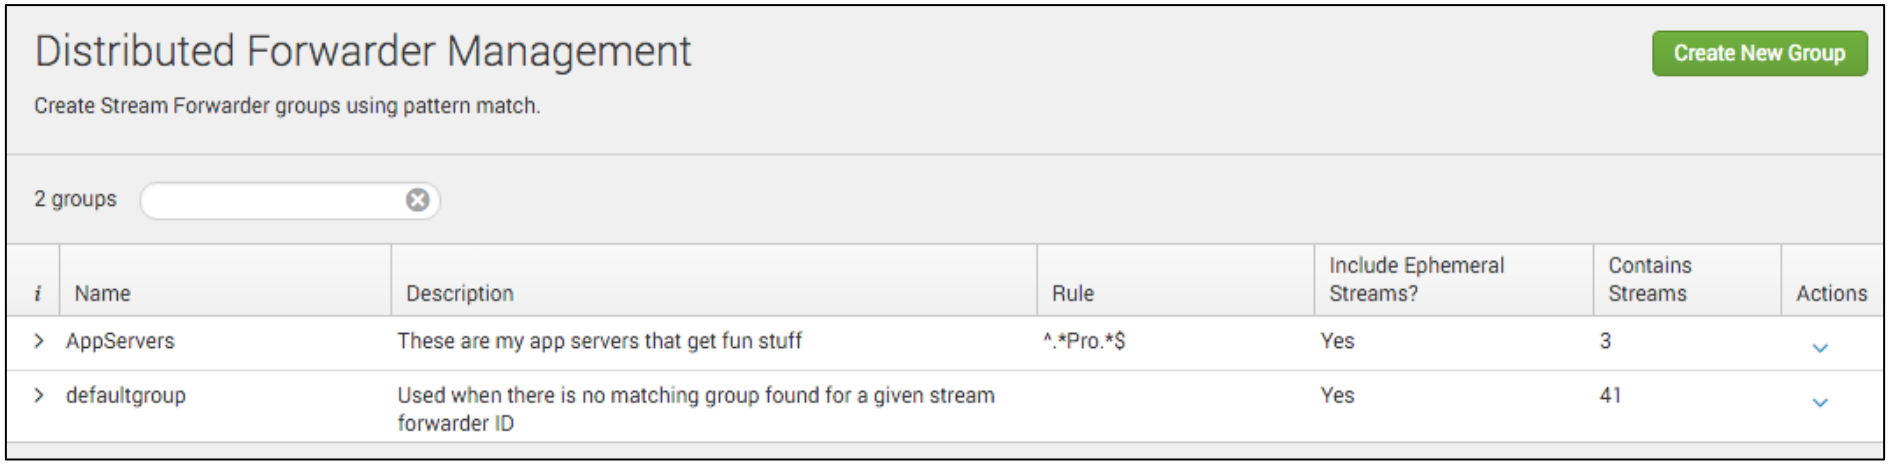

Where to use Splunk Stream

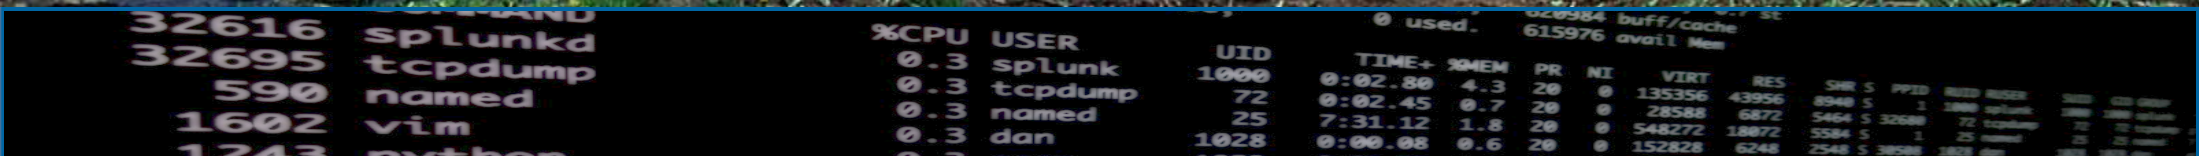

APLURA

Many Solutions, One Goal.

Dedicated Stream Forwarders

- Send data off of a switch Span or Tap
- Tools like Gigamon, Ixia, Etc.
  - You need these for really big pipes to spread the love
- Purpose built
  - Higher CPU and RAM
  - Better network cards
- Also a good option is you want to perform SSL decrypt
- Note that if you do this you will want to change some of your kernel settings (buffer sizes)
- Make sure to monitor your forwarders for thruput warnings!

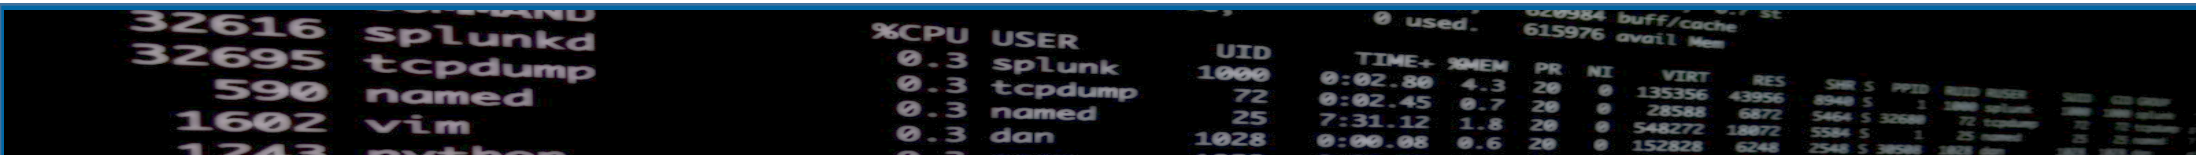

APLURA
Many Solutions, One Goal.

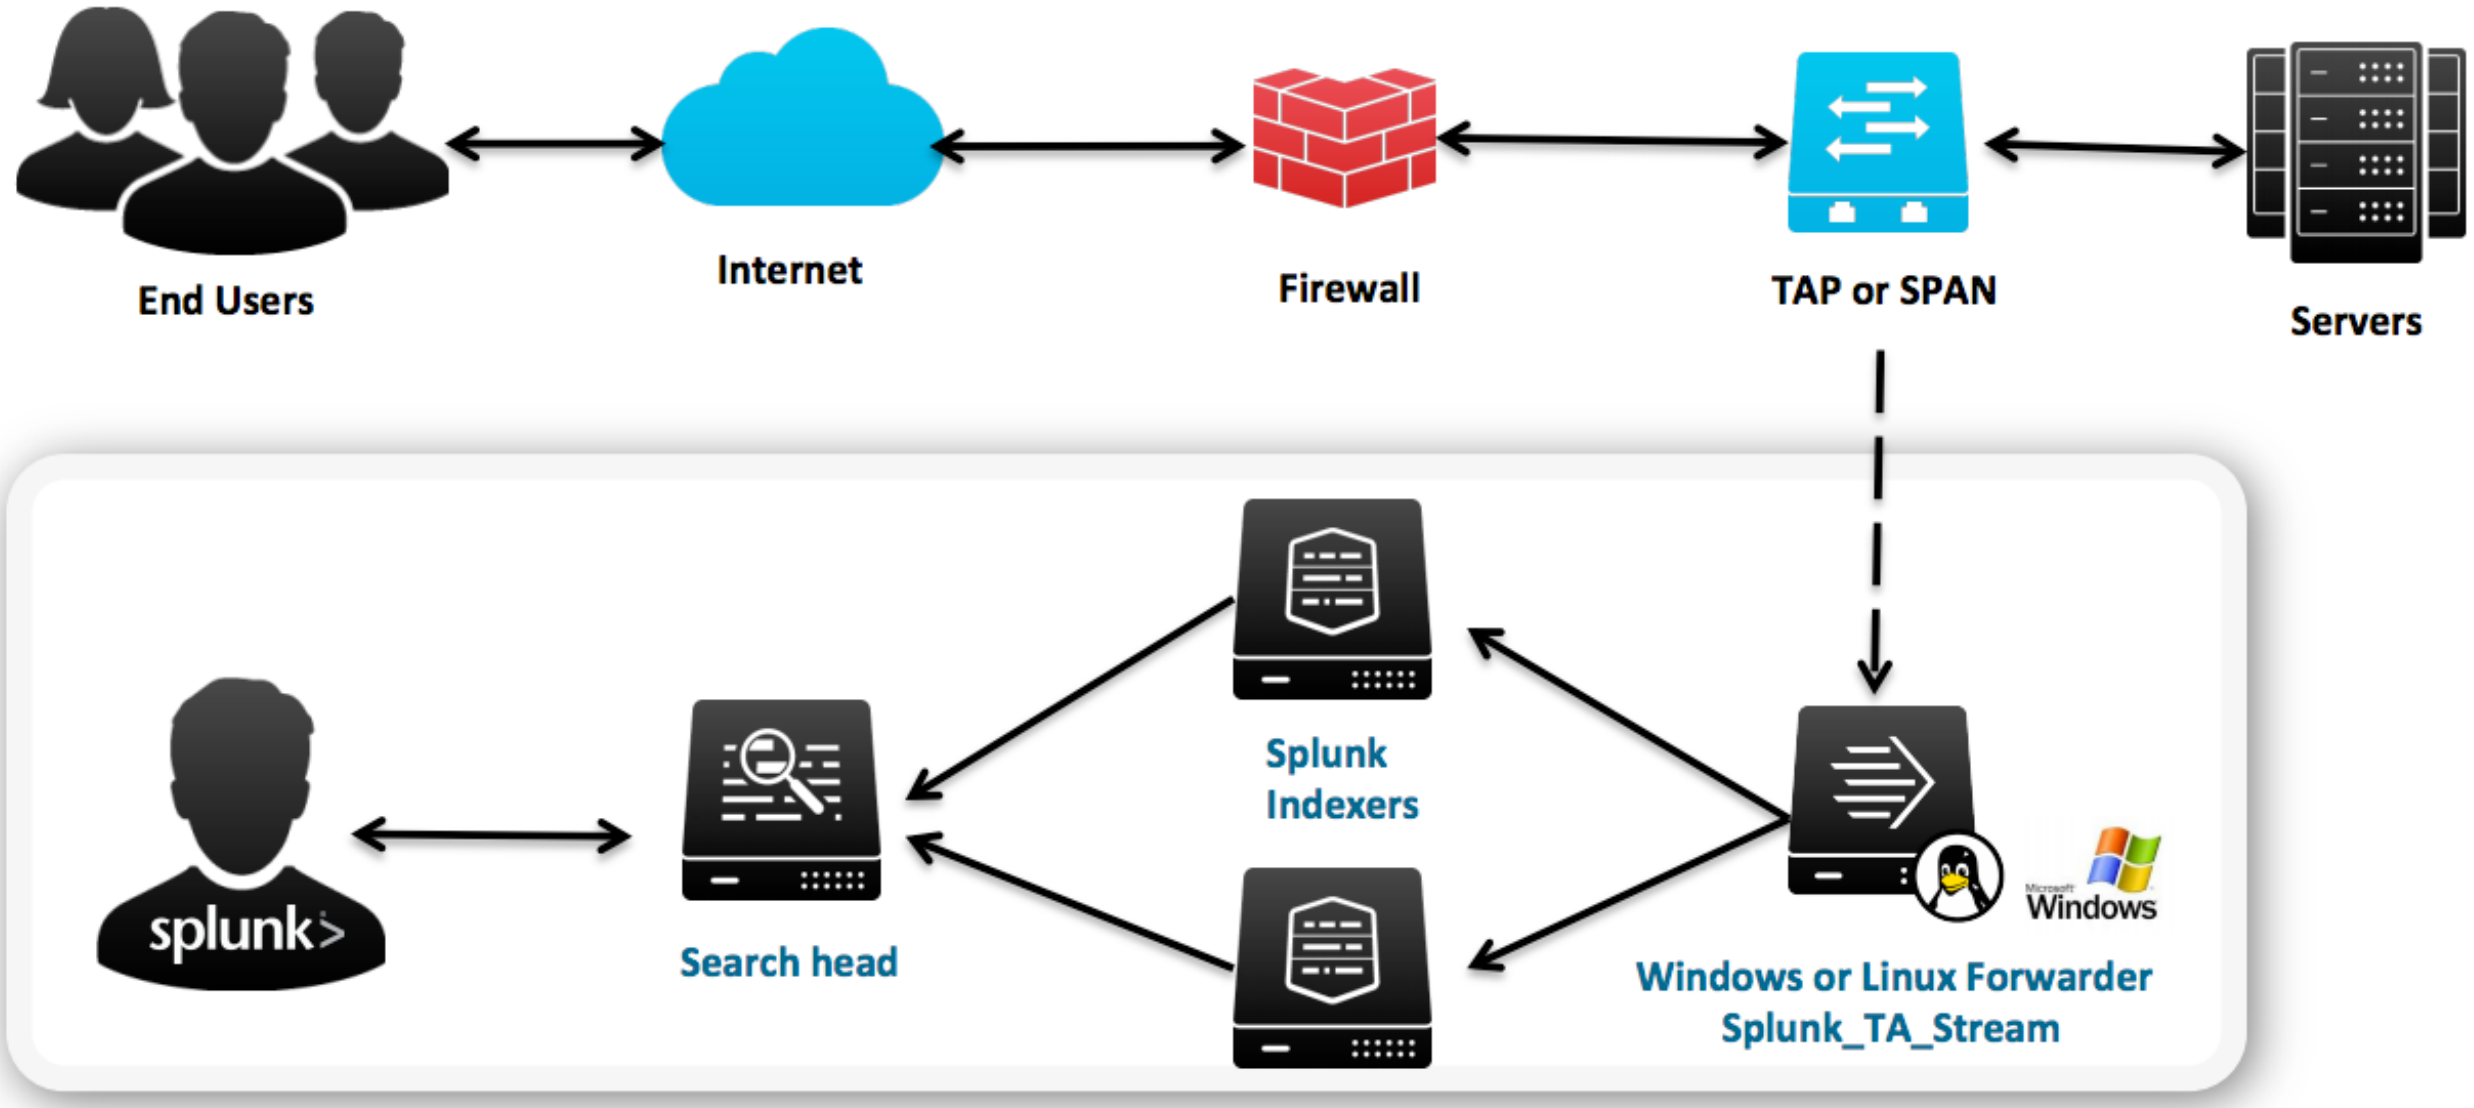
End Users
Internet
Firewall
TAP or SPAN
Servers
splunk>
Search head
Splunk
Indexers
Windows
Windows or Linux Forwarder
Splunk_TA_Stream

Deploy to the Endpoints

- Deploy directly to the systems you want to monitor
- Good for application debugging
- Nice option for Splunk ES
- Can be done from Deployment Server
- Granular control over groups
- Could mean a lot of “hand on”

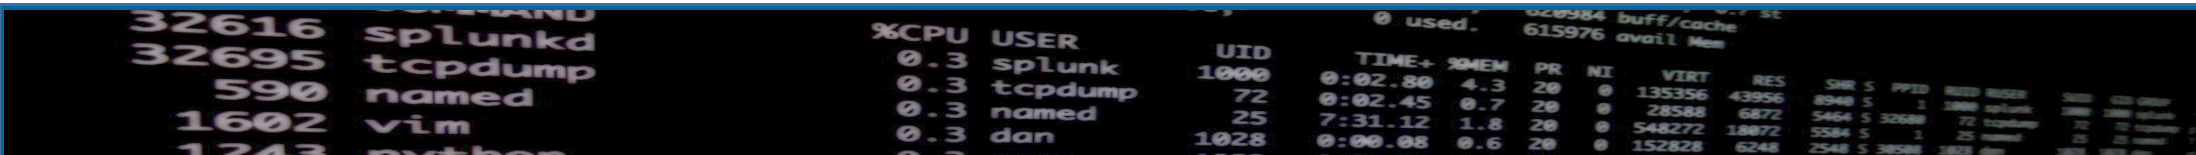

APLURA

Many Solutions, One Goal.

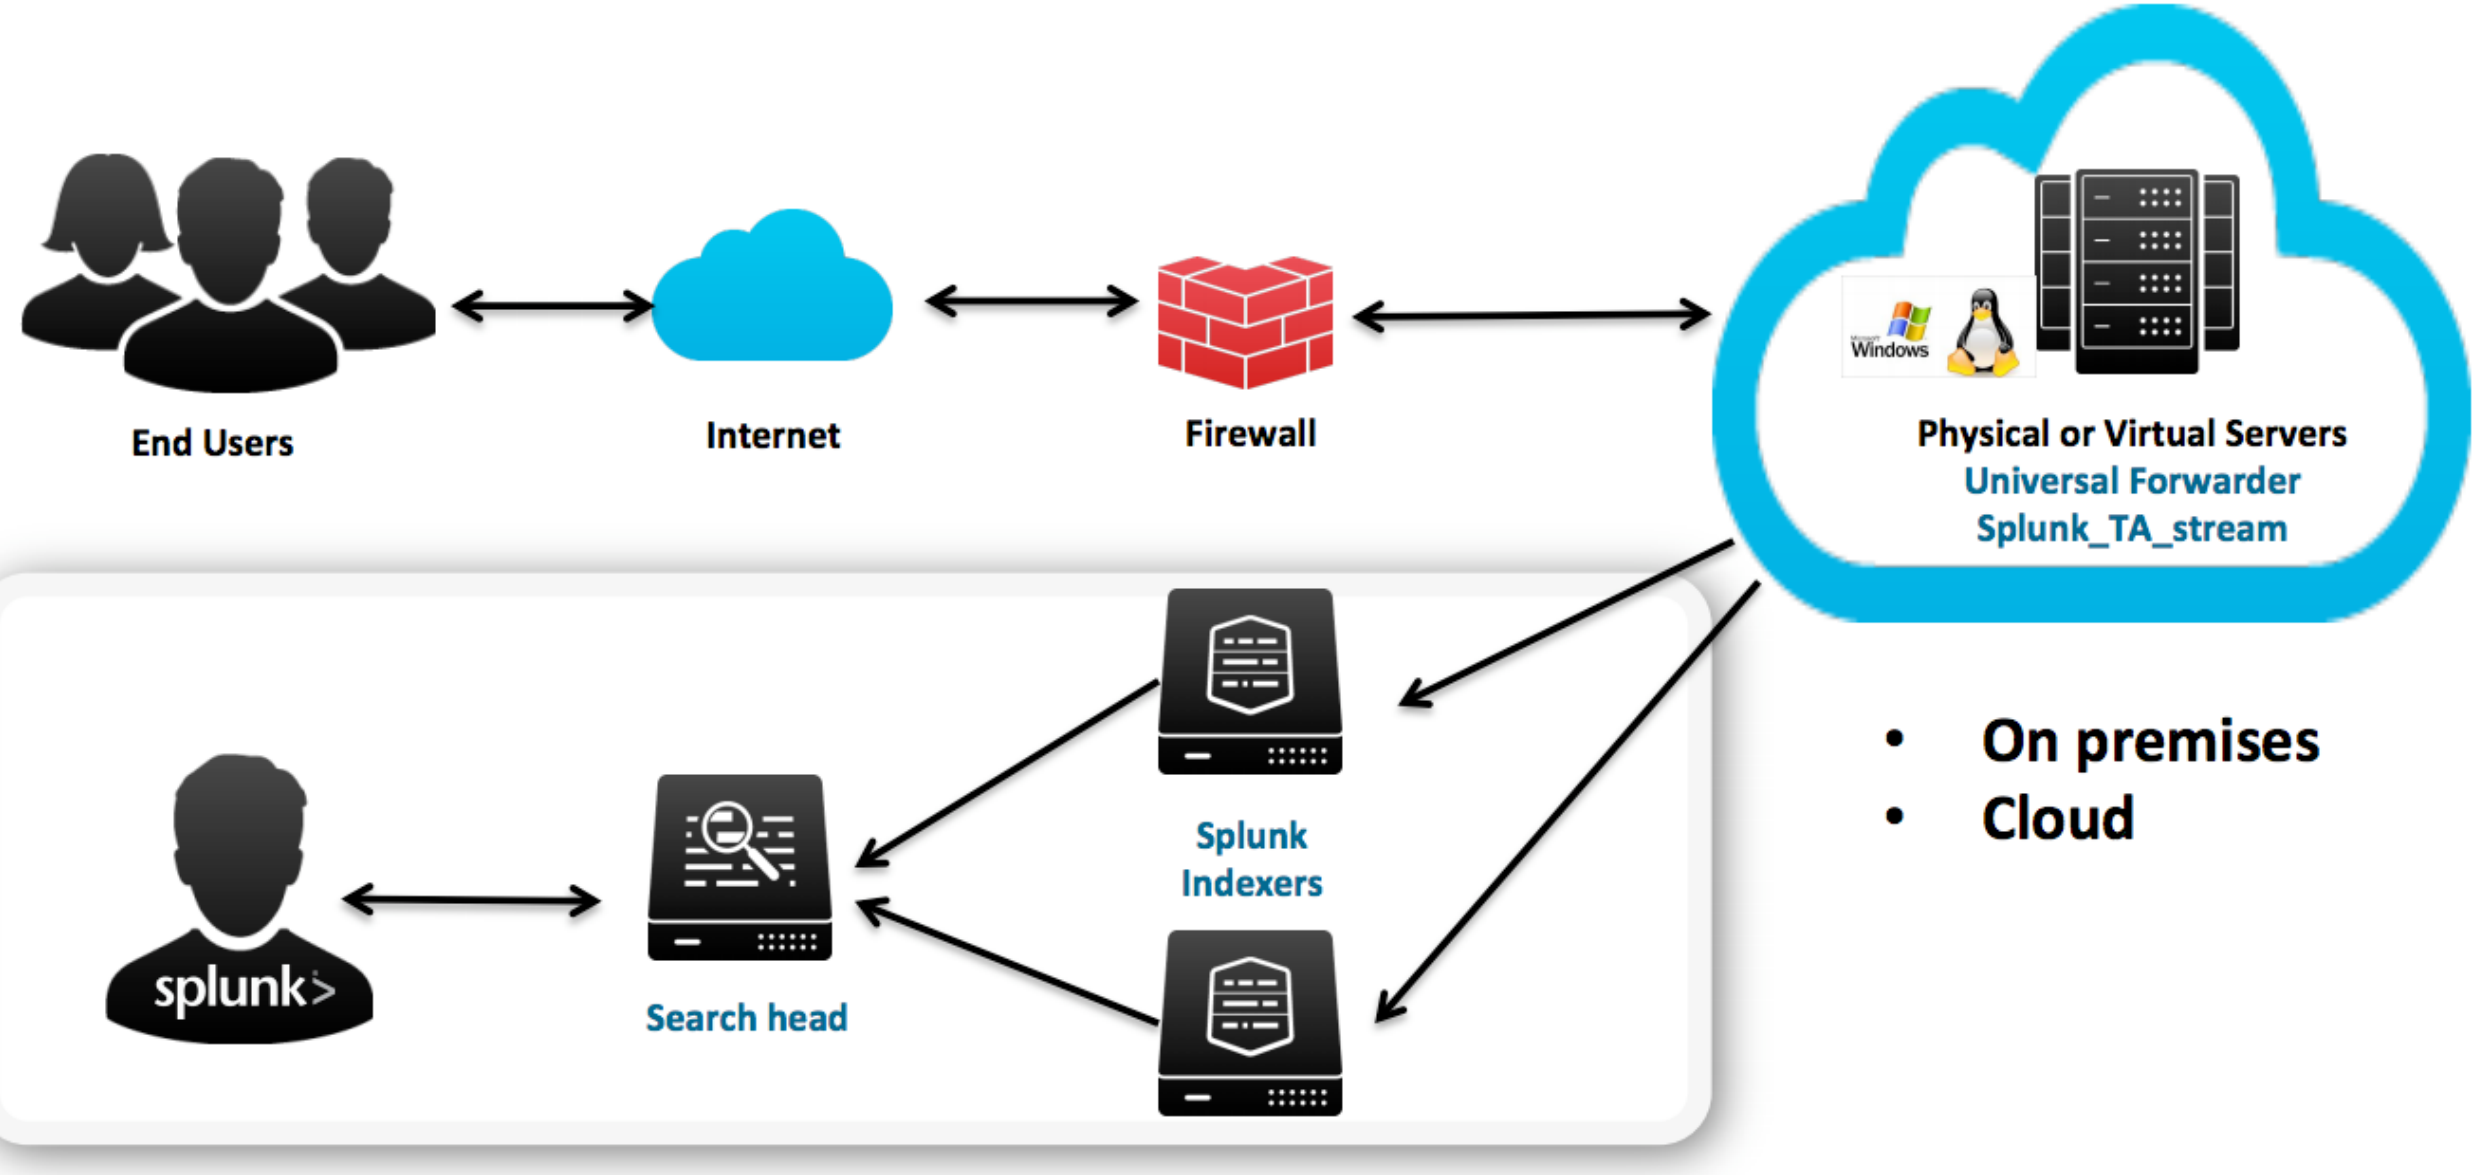

- On premises
- Cloud

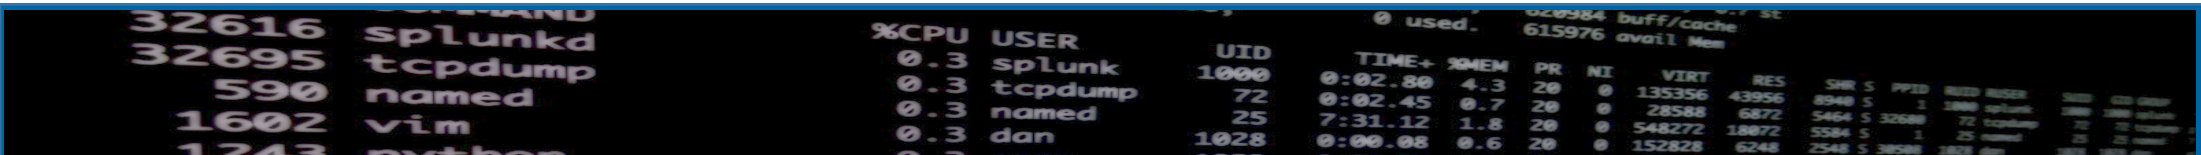

APLURA

Many Solutions, One Goal.

Deploying Splunk Stream

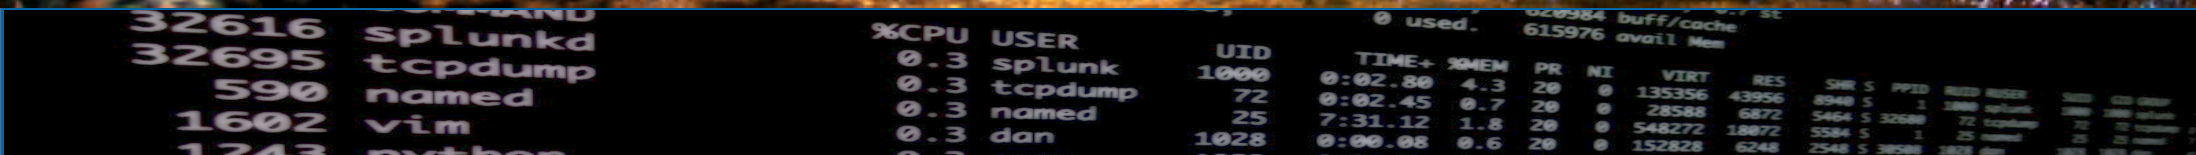

APLURA

Many Solutions, One Goal.

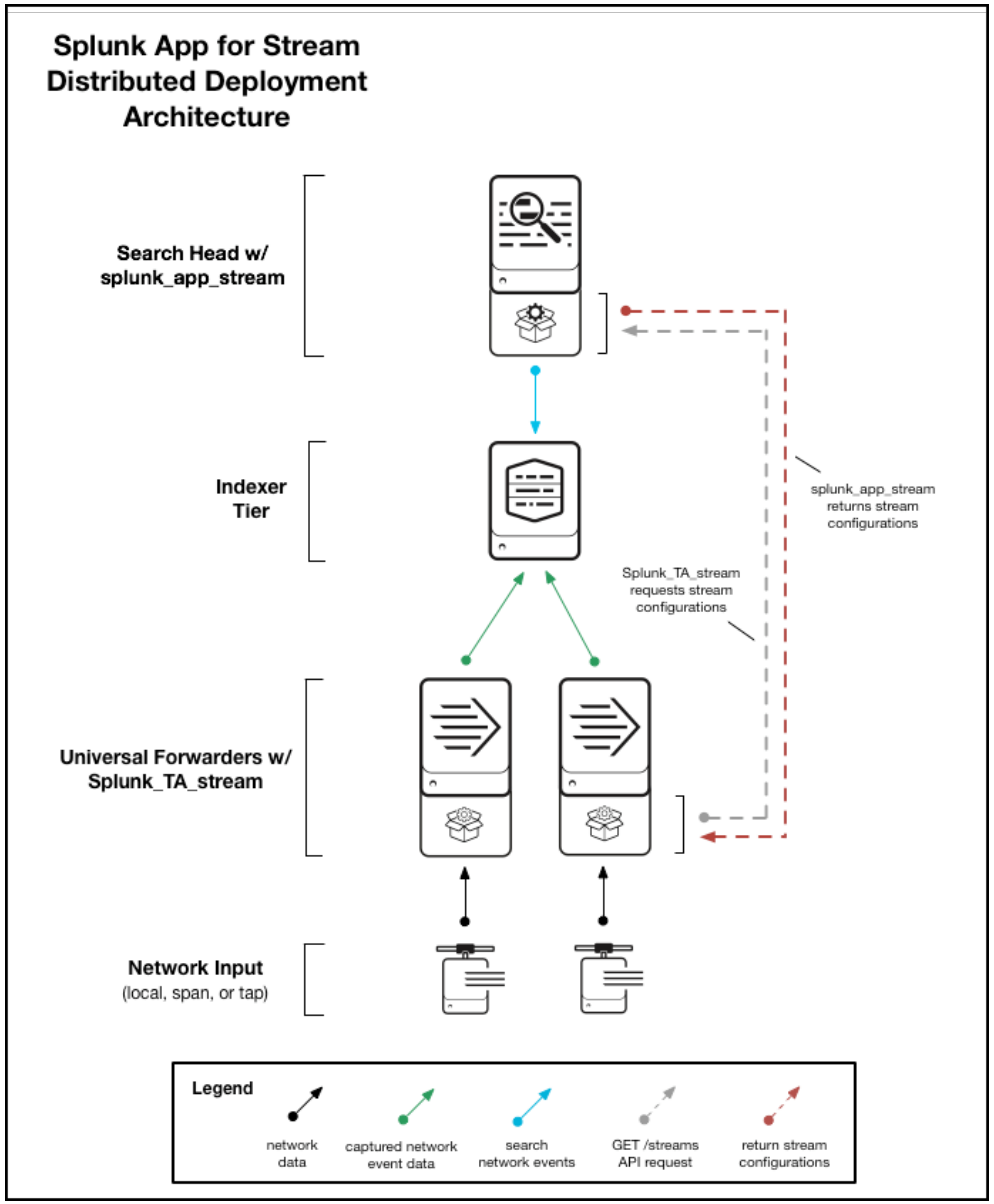

Two parts

- The Splunk App for Stream
  - Dashboards for analytics on protocols
  - Administrative panels for configuration
  - Stream Estimate (really cool, more later)
  - Goes on Search Head/Controller
- Splunk Stream Add-on
  - Binaries
  - Index-time operations (linebreaking, timestamping)
  - Goes on Indexers and Forwarders (UF or HF)

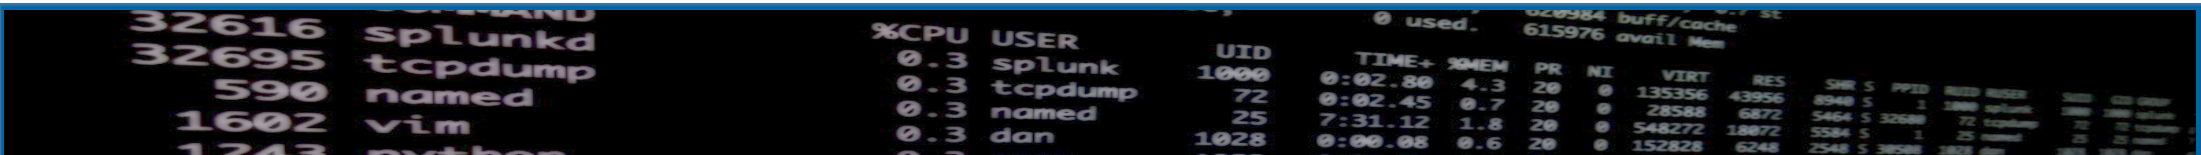

Install the Splunk App for Stream

- Can co-locate with ES
- Can co-locate with DMC
- In smaller (less than 100 forwarders) don’t use with the DS
  - Possible exhausted connections (DS and Stream poll separately)
- Installs just like any other Splunk app

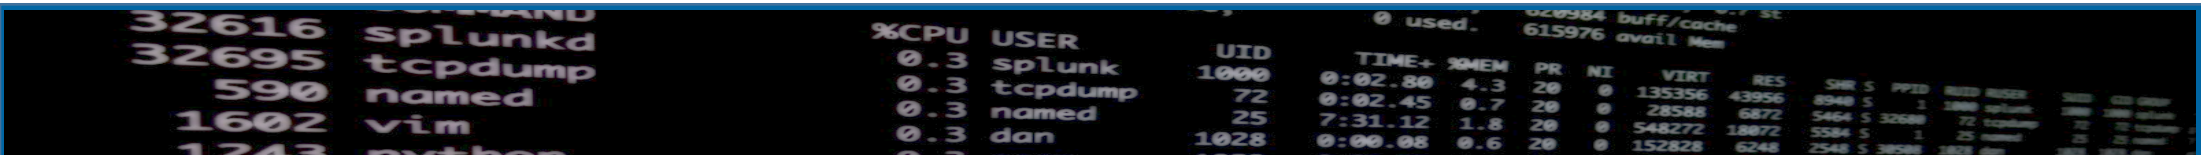

APLURA
Many Solutions, One Goal.

Harvest the Add On

- Installs to a few places
- $SPLUNK_HOME/etc/apps/Splunk_TA_stream
- $SPLUNK_HOME/etc/apps/splunk_app_stream/install/Splunk_TA_stream
- $SPLUNK_HOME/etc/deployment-apps/Splunk_TA_stream
  - Will create the local inputs.conf with the app server location

```
Davids-MacBook-Pro-2:local root# cat inputs.conf
[streamfwd://streamfwd]
splunk_stream_app_location = https://Davids-MacBook-Pro-2.local:8000/en-us/custom/splunk_app_stream/
stream_forwarder_id =
disabled = 0
```

* Skip this is your SH is your DS

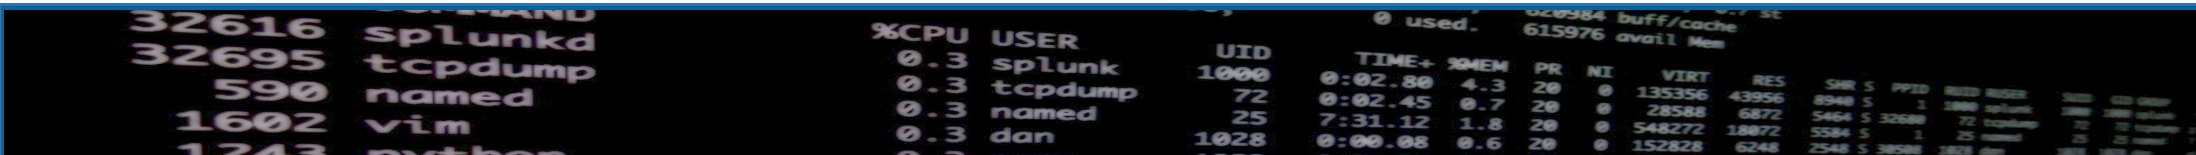

APLURA
Many Solutions, One Goal.

Make sure your forwarders can talk back

- Your forwarders will need to be able to talk to the SH with splunk_app_stream installed
- The port is the same as the GUI for your SH

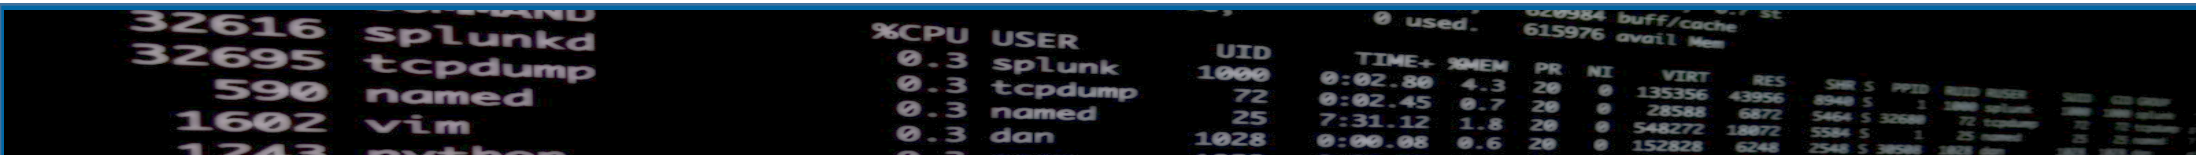

Configure your forwarders

- Don’t have to be root on Linux
  - Use the included setuid.sh script
- Must be local admin or local system on Windows
- On UFs you should monitor your thruput limits

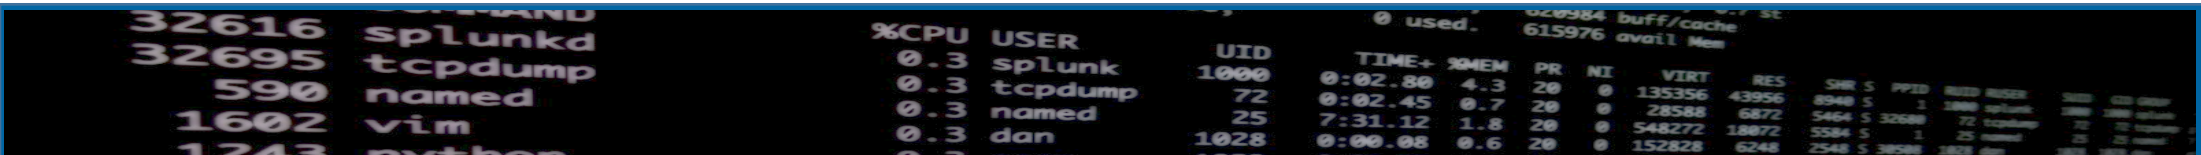

Inputs.conf

- Remember that the inputs.conf is layerable
- Just like other Splunk configs
- Doesn’t have to be in the Splunk_TA_stream
- On the DS you can deploy two apps, one with the input to point back to the splunk_app_stream
- Then also deploy the Splunk_TA_stream

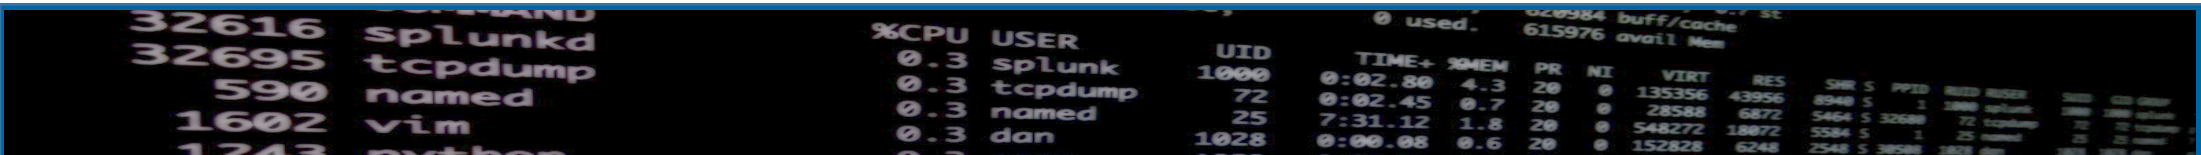

Many Solutions, One Goal.

Configure your streams

- The defaults may send more fields than you need
- Can tell forwarders which parts of the data you want
- You can have different configs for different groups!

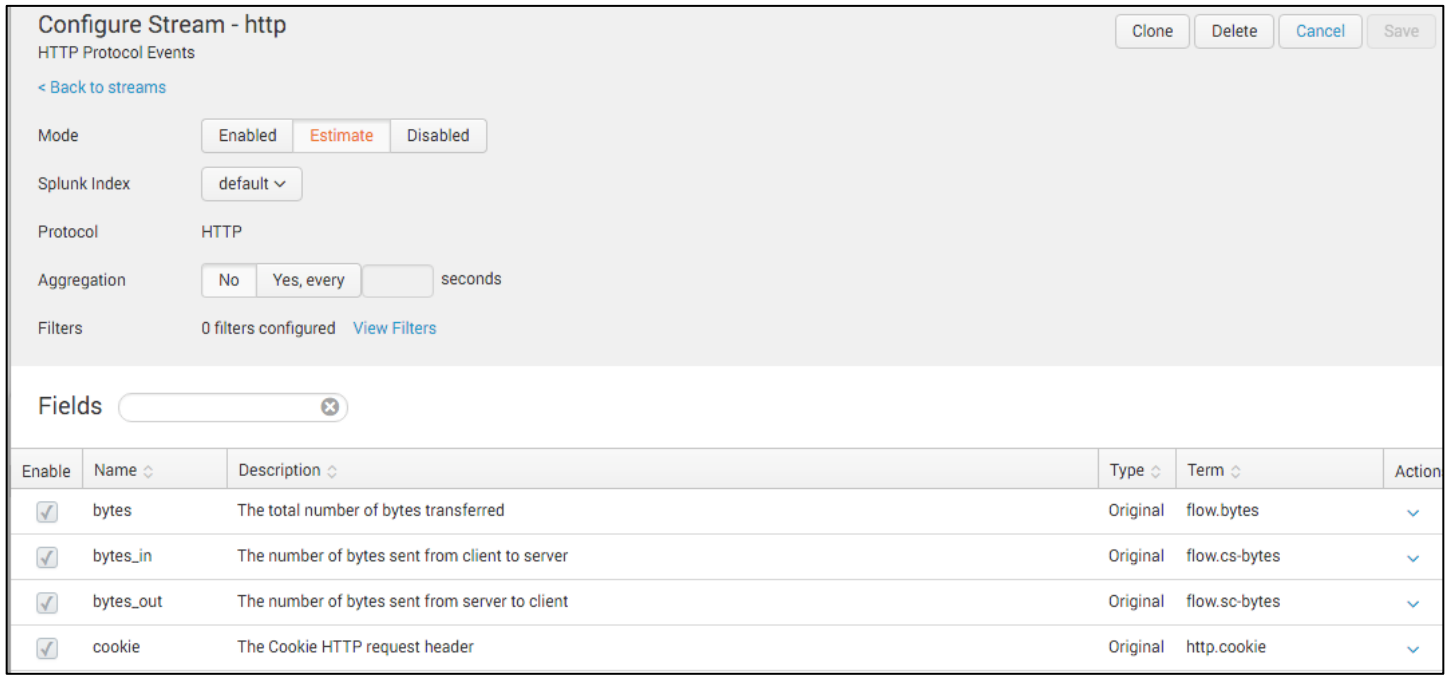

Configure your forwarder groups

- Uses good ol' regex
- Lets you say ahead of time if Ephemeral Streams should be allowed

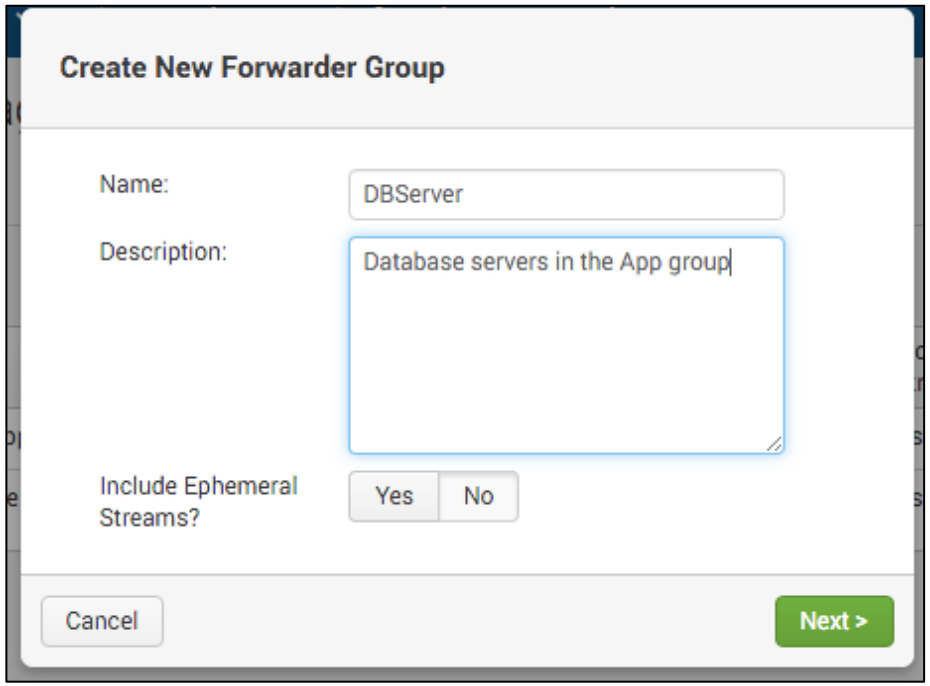

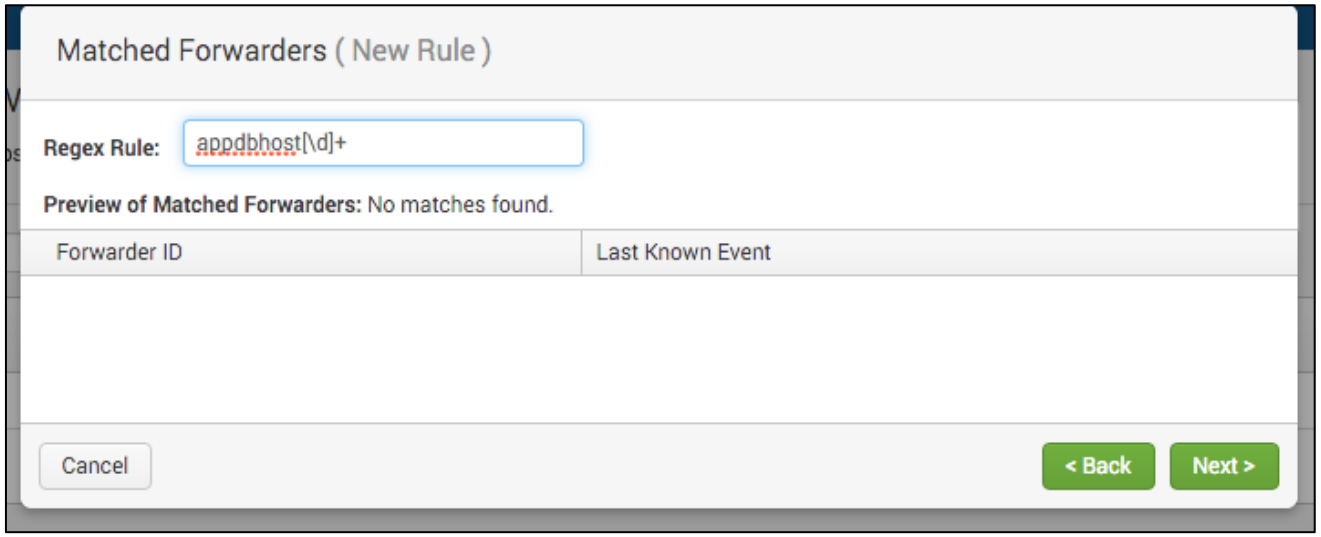

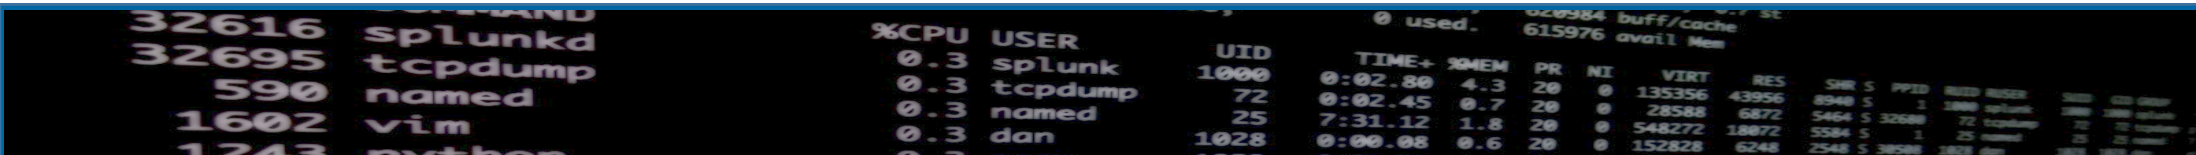

APLURA
Many Solutions, One Goal.

Gotcha with Groups

- Just regex on the Stream forwarder ID (not IP, hostname)
- This is configured in an XML file
- Messy
- The “defaultgroup” forwarder group for all unmatched hosts will gather ALL THE THINGS

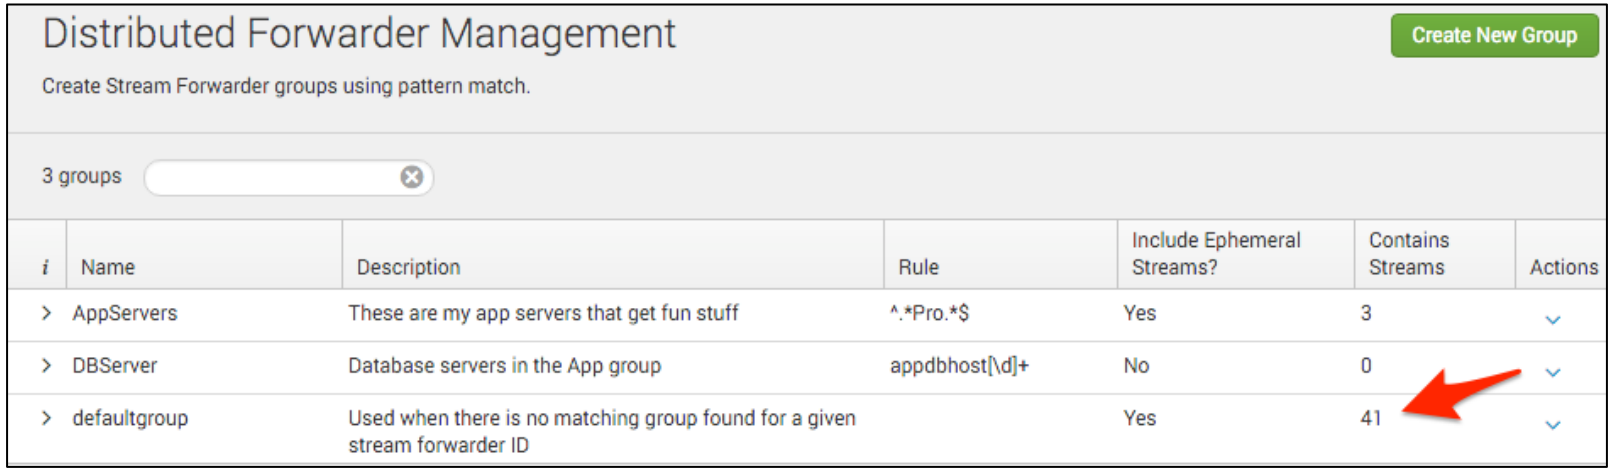

Wait for data to flow in

- That’s pretty much it!
- Docs make it look a lot harder

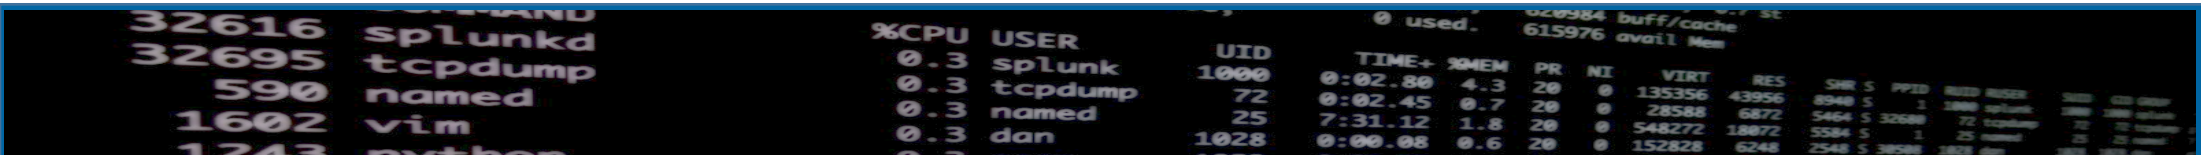

Questions?

APLURA

Many Solutions, One Goal.

Credits

- Thanks to the Baltimore Area Splunk User Group
- Cover Slide: Upper Swallow Falls in Oakland, MD, Chris Flees, http://fineartamerica.com/profiles/chris-flees.html?tab=artwork&page=7
- Slide 3: Potomac River in Maryland, Terry J. Adams, http://www.fhwa.dot.gov/byways/byways/60807/photos
- Slide 7: Timanus Mill on the Jones Falls in Baltimore, “Monument City”, http://www.panoramio.com/photo/57148558
- Slide 8: “Missing Homework Log” by “Red Beetle RB”. https://www.teacherspayteachers.com/Product/Missing-Homework-Log-4112
- Slide 9: Rotton log, National Wildlife Foundation, https://www.nwf.org/kids/family-fun/outdoor-activities/investigate-a-rotten-log.aspx
- Slide 10: The Simpsons, http://i.imgur.com/91sn32Q.jpg?fb
- Slide 11: Bro Network Security Monitor, https://www.bro.org/
- Slide 17: Ian Adams Photography, http://ianadamsphotography.com/news/galleries/bridges/
- Slides 19 and 21: Splunk Conf 2015, “Splunk App for Stream Deployments in the Real World: Enhance Operational Intelligence Across Application Delivery, IT Ops, Security and More”, http://conf.splunk.com/session/2015/conf2015_SUdovicic_CChing_MDickey_Splunk_SplunkEntWhatsNew_StreamDeploymentsInTheReal.pdf
- Slide 22: Gunpowder Falls in Baltimore County, MD, http://hdrcreme.com/photos/1818-gunpowder-falls
- Slide 23: Splunk Docs, http://docs.splunk.com/Documentation/StreamApp/latest/DeployStreamApp/DeploymentArchitecture
- Slide 34: Youghiogheny River at Friendsville, MD by Joe Dawson, https://www.flickr.com/photos/jmd41280/5066756138

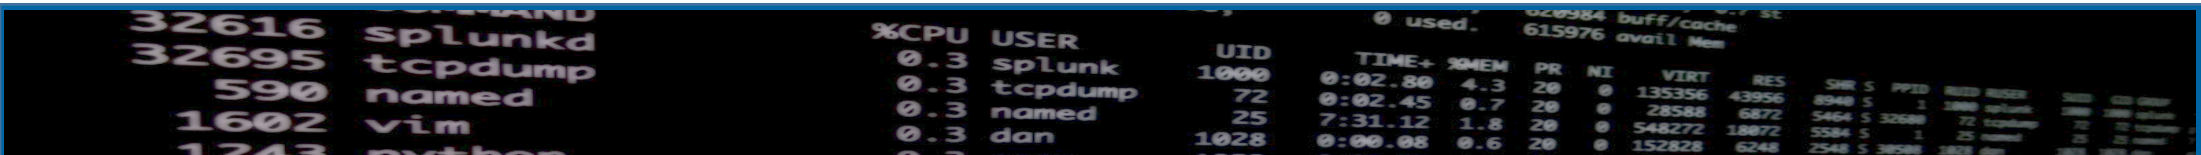

APLURA
Many Solutions, One Goal.