Splunk Architecture Workshop

We Got Splunk
Its Installed
Cool
Now What?

APLURA

Who Am I?

- Sean Wilkerson, Partner/Consultant, Aplura

Who Am I?

- Sean Wilkerson, Partner/Consultant, Aplura
- ~15 Years of Network --> Systems --> **InfoSec**
- A Decade+ of Federal Log-Management
  - Half Spent Deploy/Manage FOSS/SIM/SIEM
  - Half Spent Deploy/Manage FOSS/Splunk
- 2 Recent SANS Log Mgmt Summits
- Splunk Pro Serv Partner Since 2008
- Splunk makes me happy

APLURA

Who Are You?

- You Know Splunk is Right for You
- You Know Key Splunk Concepts
- You Require Purpose-Built Architecture
- You Have Scale
  - Big Data
  - Lots of Users
  - Lots of Data-types
  - Highly Distributed

Splunk PS Work

Major Activities

- App Integration
- App Development
- Enhancements
- Repairs

Architecture

- Migrations
- Merging/Breakout
- Upgrades
- DR/HA
- Build for Speed/Size

But Isn't Splunk EZ?

But Isn't Splunk EZ?

- Yes...and no, err, let me explain

But Isn't Splunk EZ?

- Yes...and no, err, let me explain
- Single-box can be Grandma-Easy

But Isn't Splunk EZ?

- Yes...and no, err, let me explain
- Single-box can be Grandma-Easy
- PT: Enterprise Needs → Enterprise Plan

But Isn't Splunk EZ?

- Yes...and no, err, let me explain
- Single-box can be Grandma-Easy
- PT: Enterprise Needs → Enterprise Plan
- Compared with SIM/SIEM Competitors:

But Isn't Splunk EZ?

- Yes...and no, err, let me explain
- Single-box can be Grandma-Easy
- PT: Enterprise Needs → Enterprise Plan
- Compared with SIM/SIEM Competitors:
  - Splunk is a walk in the park
  - Incredibly more flexible
  - More transparency

But Isn't Splunk EZ?

- Yes...and no, err, let me explain
- Single-box can be Grandma-Easy
- PT: Enterprise Needs → Enterprise Plan
- Compared with SIM/SIEM Competitors:
  - Splunk is a walk in the park
  - Incredibly more flexible
  - More transparency

Note: PT = ProTip … These are keepers

Content Available Now!

- Talk: aplura.com/splunklivedc
- BP: aplura.com/splunkbp

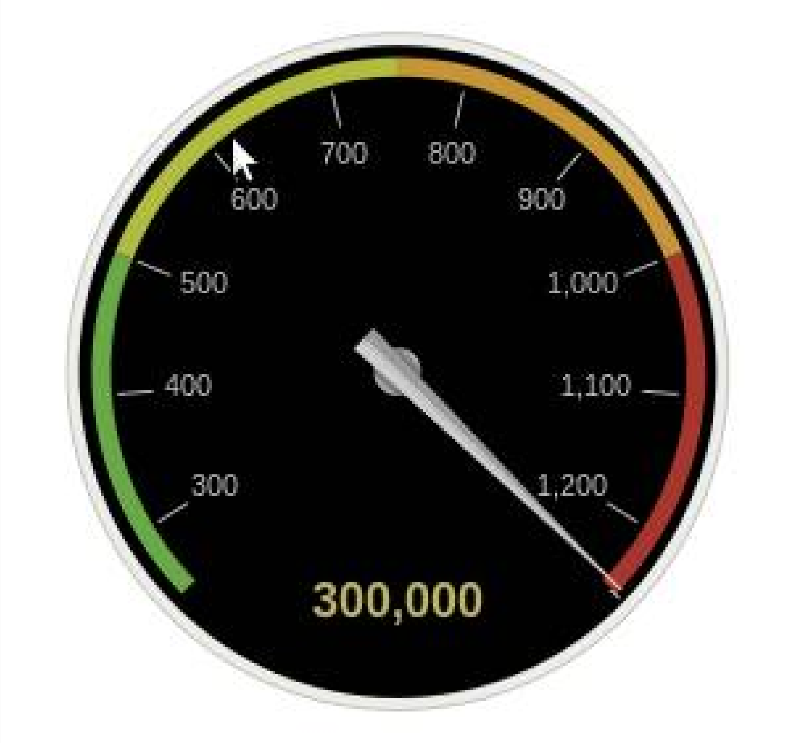
300
400
500
600
700
800
900
1,000
1,100
1,200
300,000

Sizing Your Deployment

- See Previous Talk(s)
- Key Elements
  - Architecture Supports Volume (Data and Usage)
  - Carefully Plan: Systems, Storage, and Placement
  - Administration

Reference Architecture in <5 Min

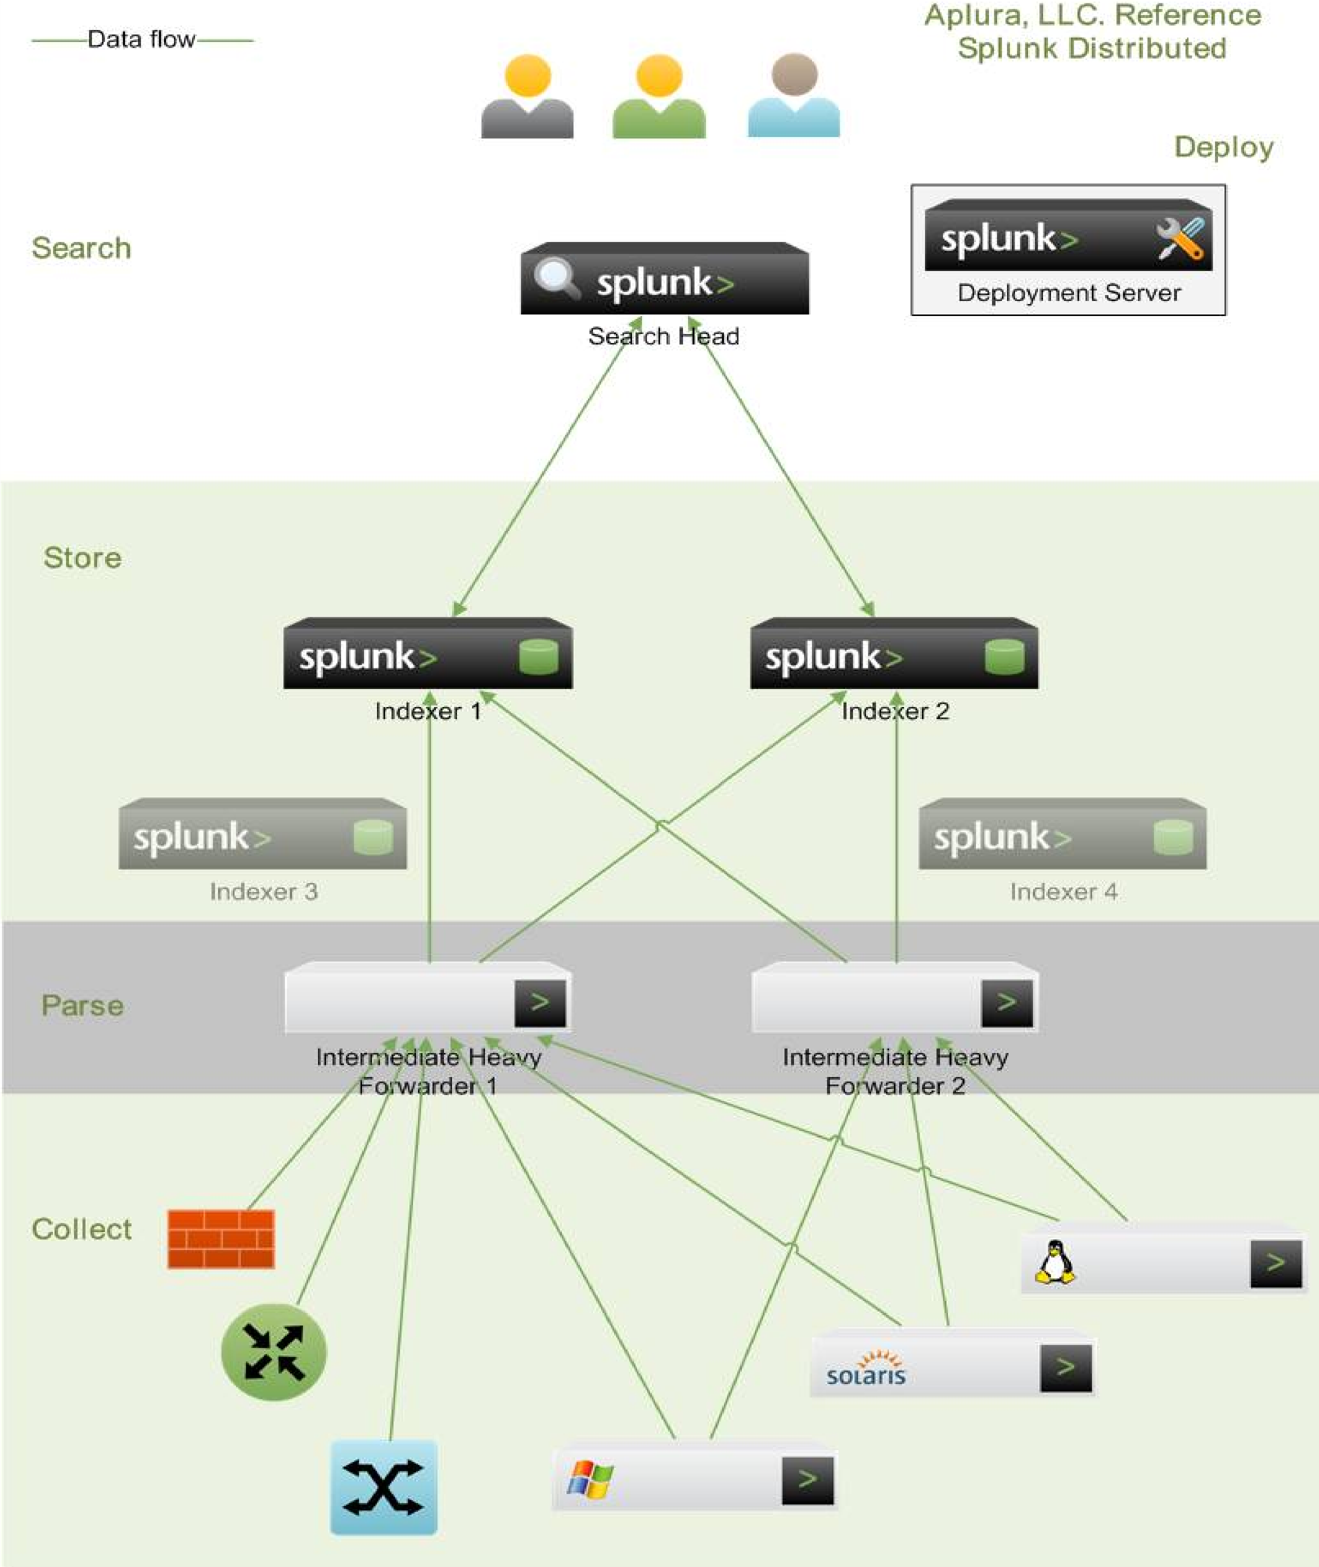

Aplura, LLC. Reference
Splunk Distributed

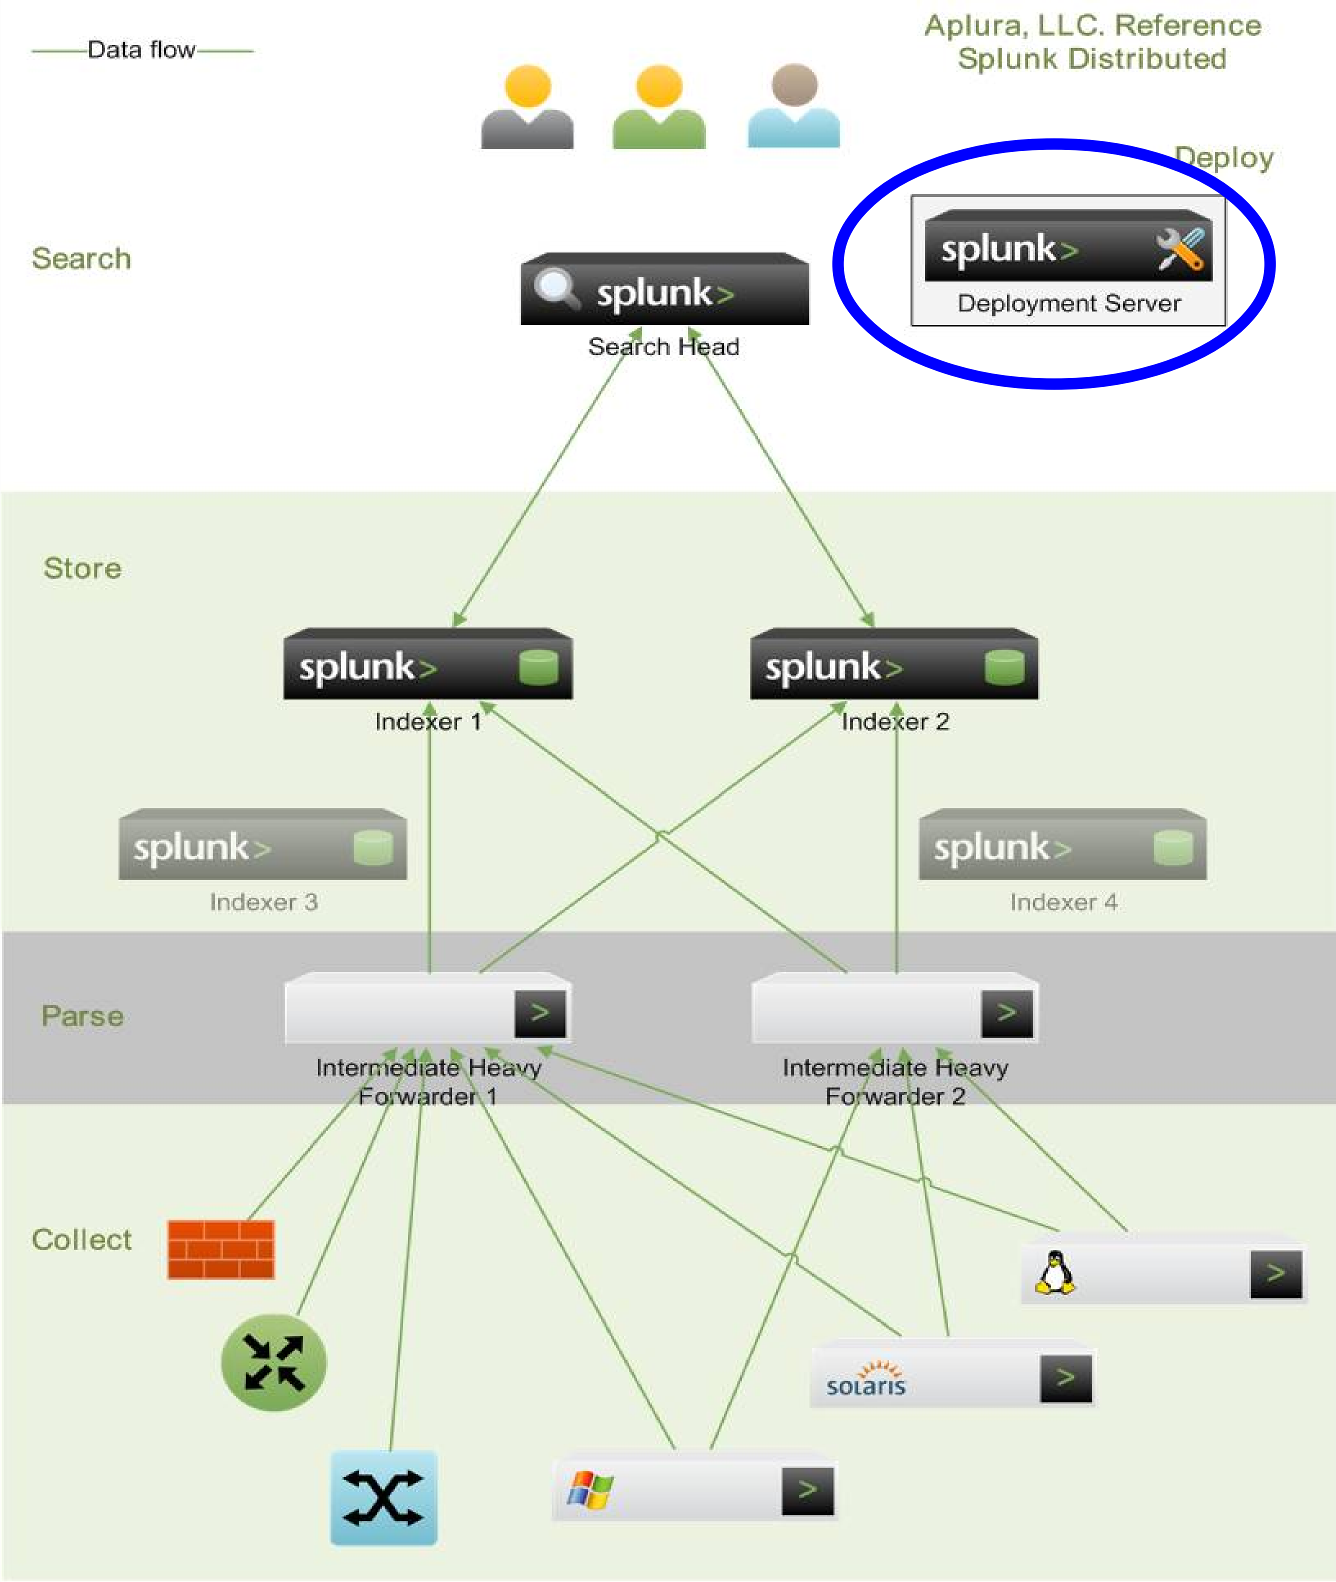

Aplura, LLC. Reference Splunk Distributed

Deployment Server (DS):

- Manages All Splunk Configs via Apps
- Platform Independent
- Ensures Uniformity
- Allows for Rapid Changes

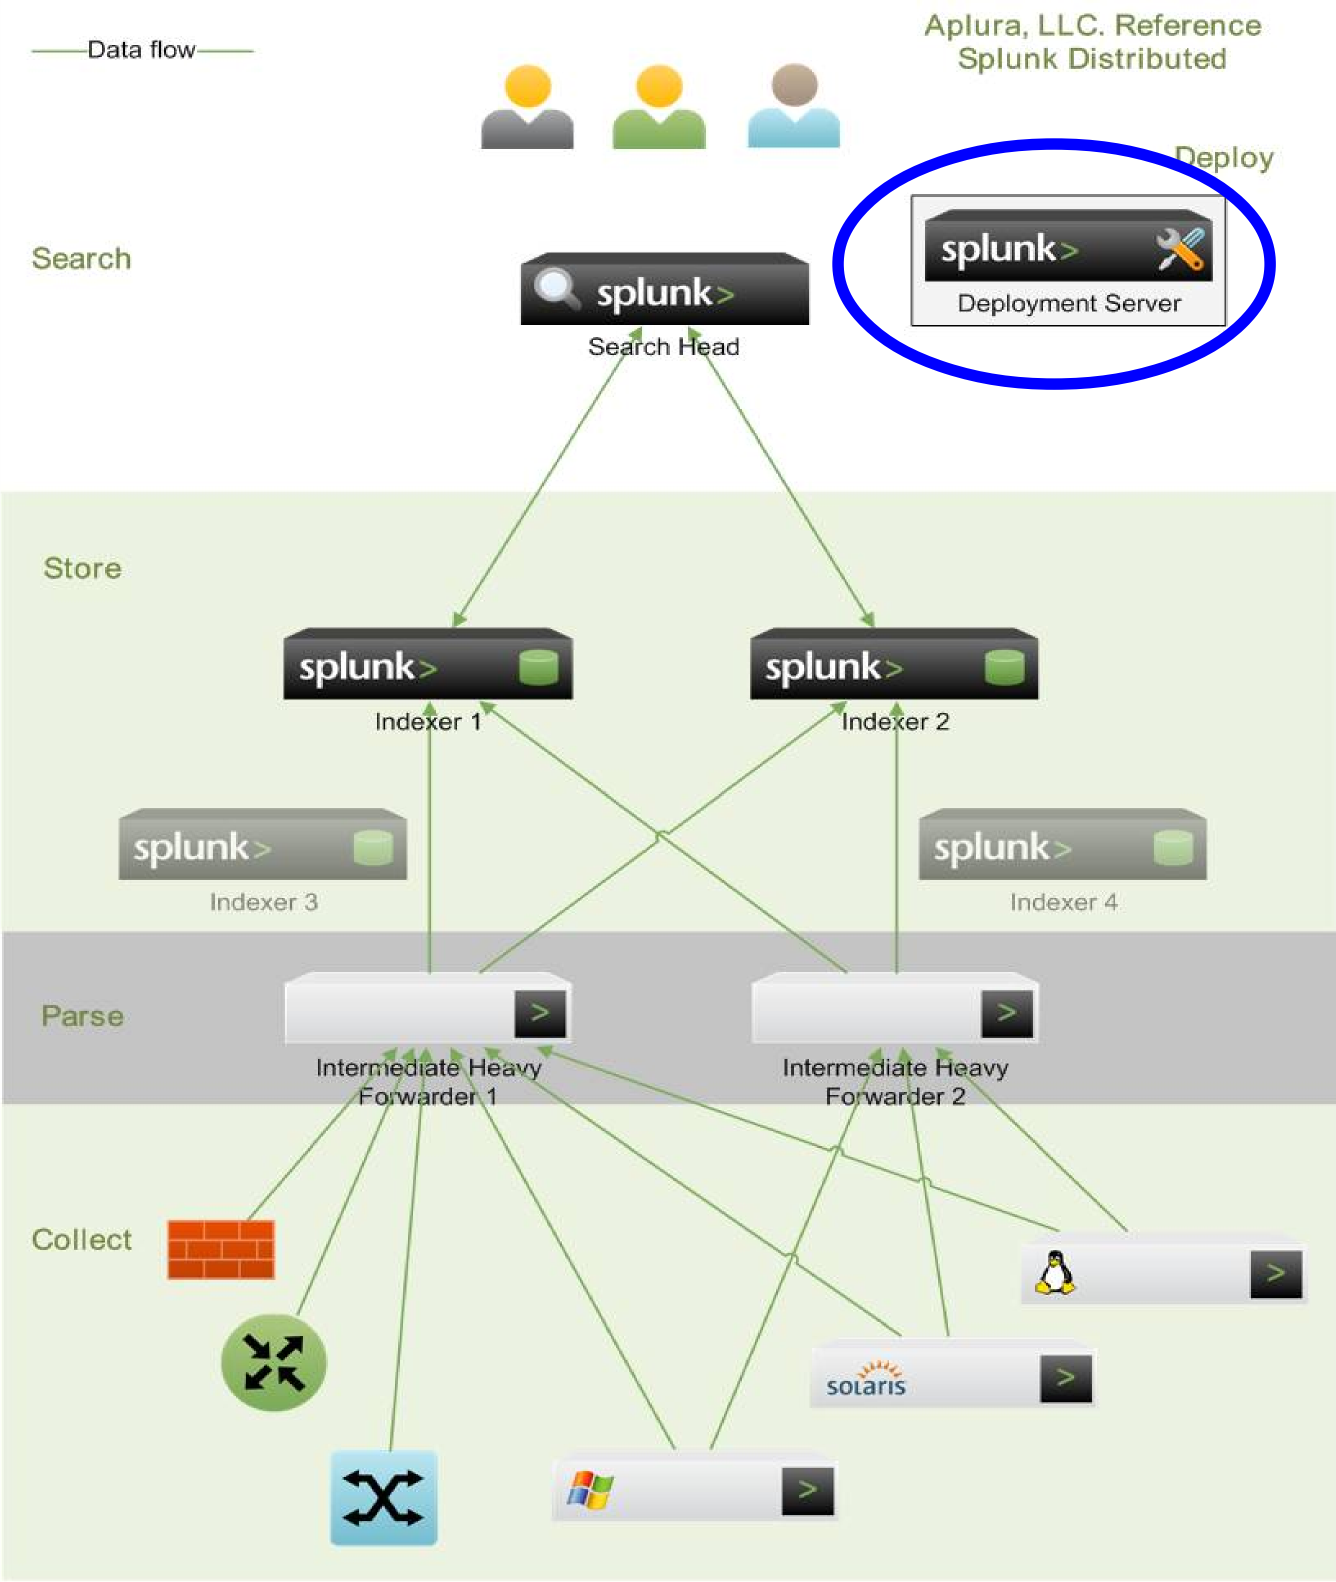

Aplura, LLC. Reference Splunk Distributed

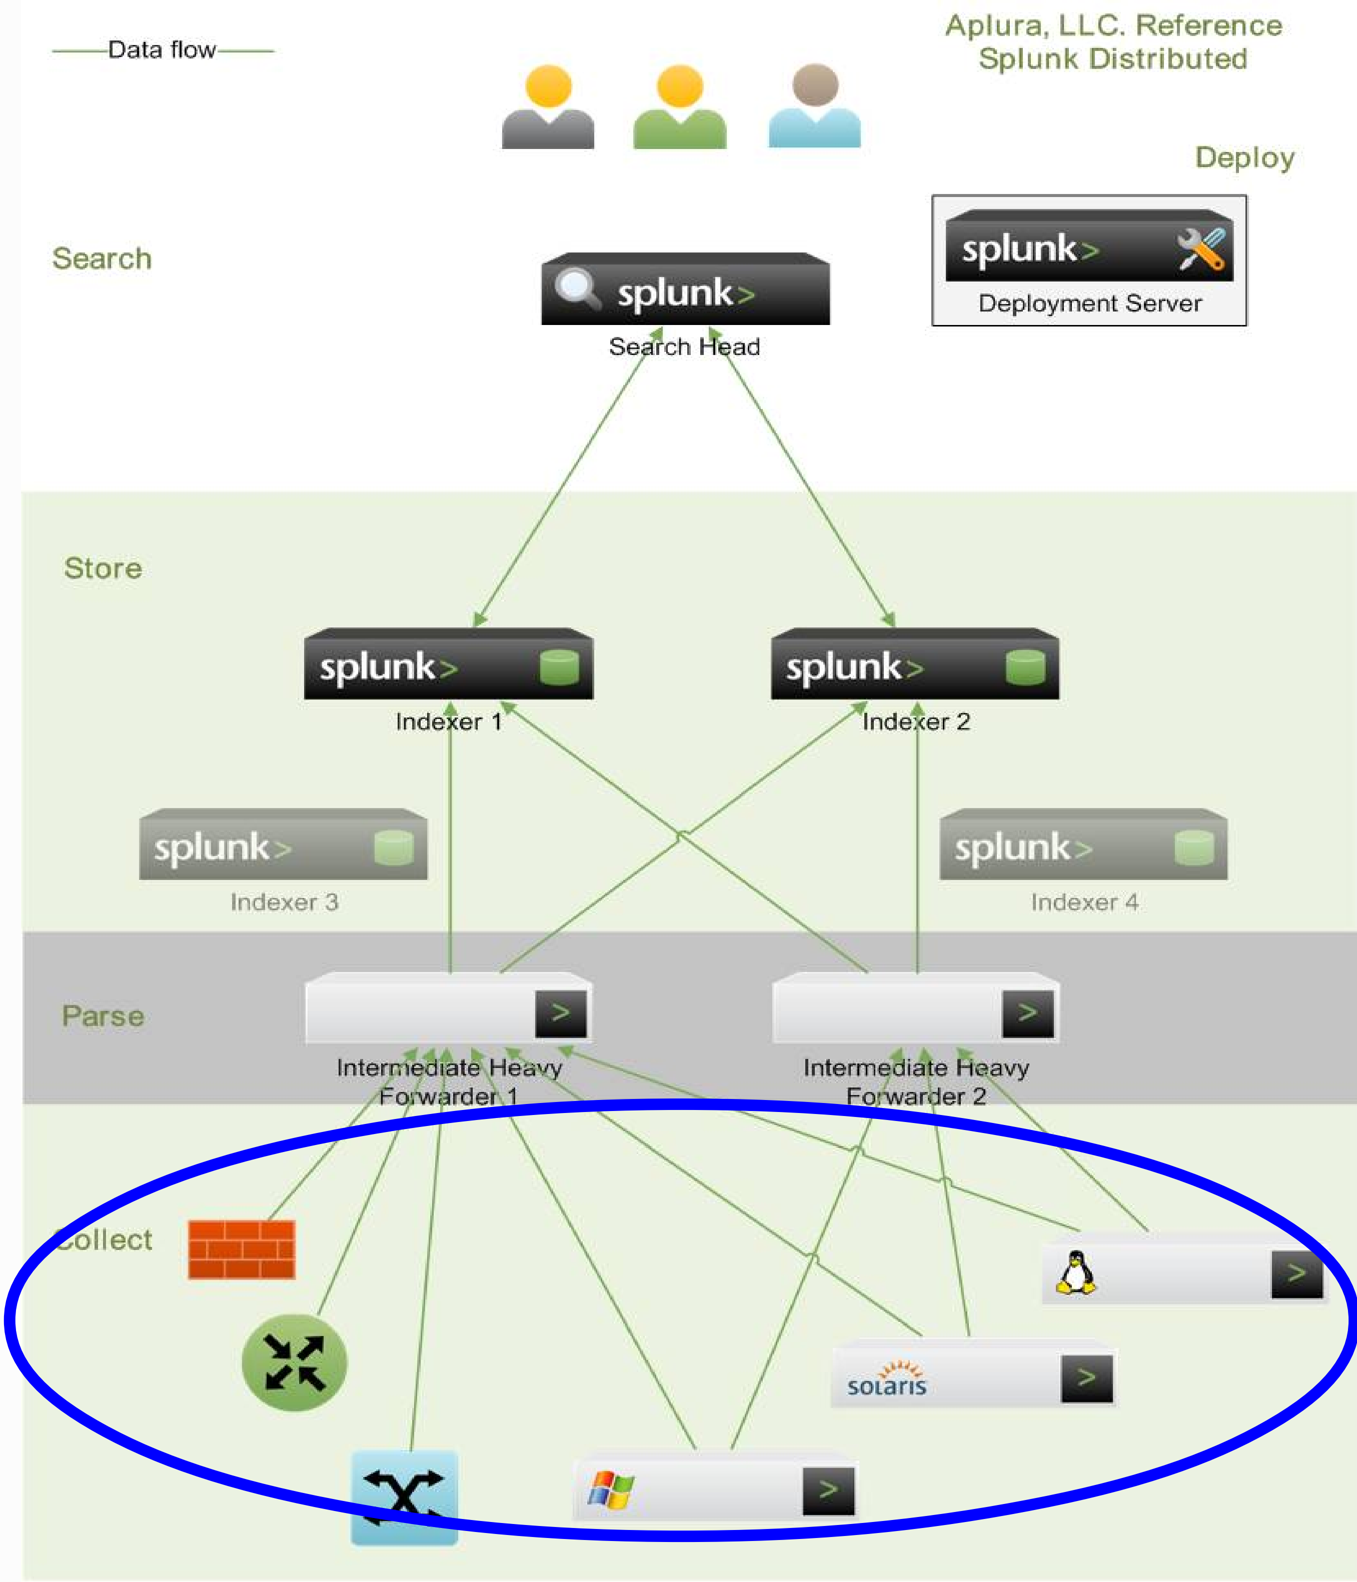

Data flow

Deploy

Deployment Server

Search

Search Head

Store

Indexer 1

Indexer 2

Indexer 3

Indexer 4

Parse

Intermediate Heavy Forwarder 1

Intermediate Heavy Forwarder 2

Collect

The Collection Tier could include:

- Syslog Receipt
- File Monitor
- DB Monitor
- Script Execution

Consider:

- Reliability
- Management
- Encryption

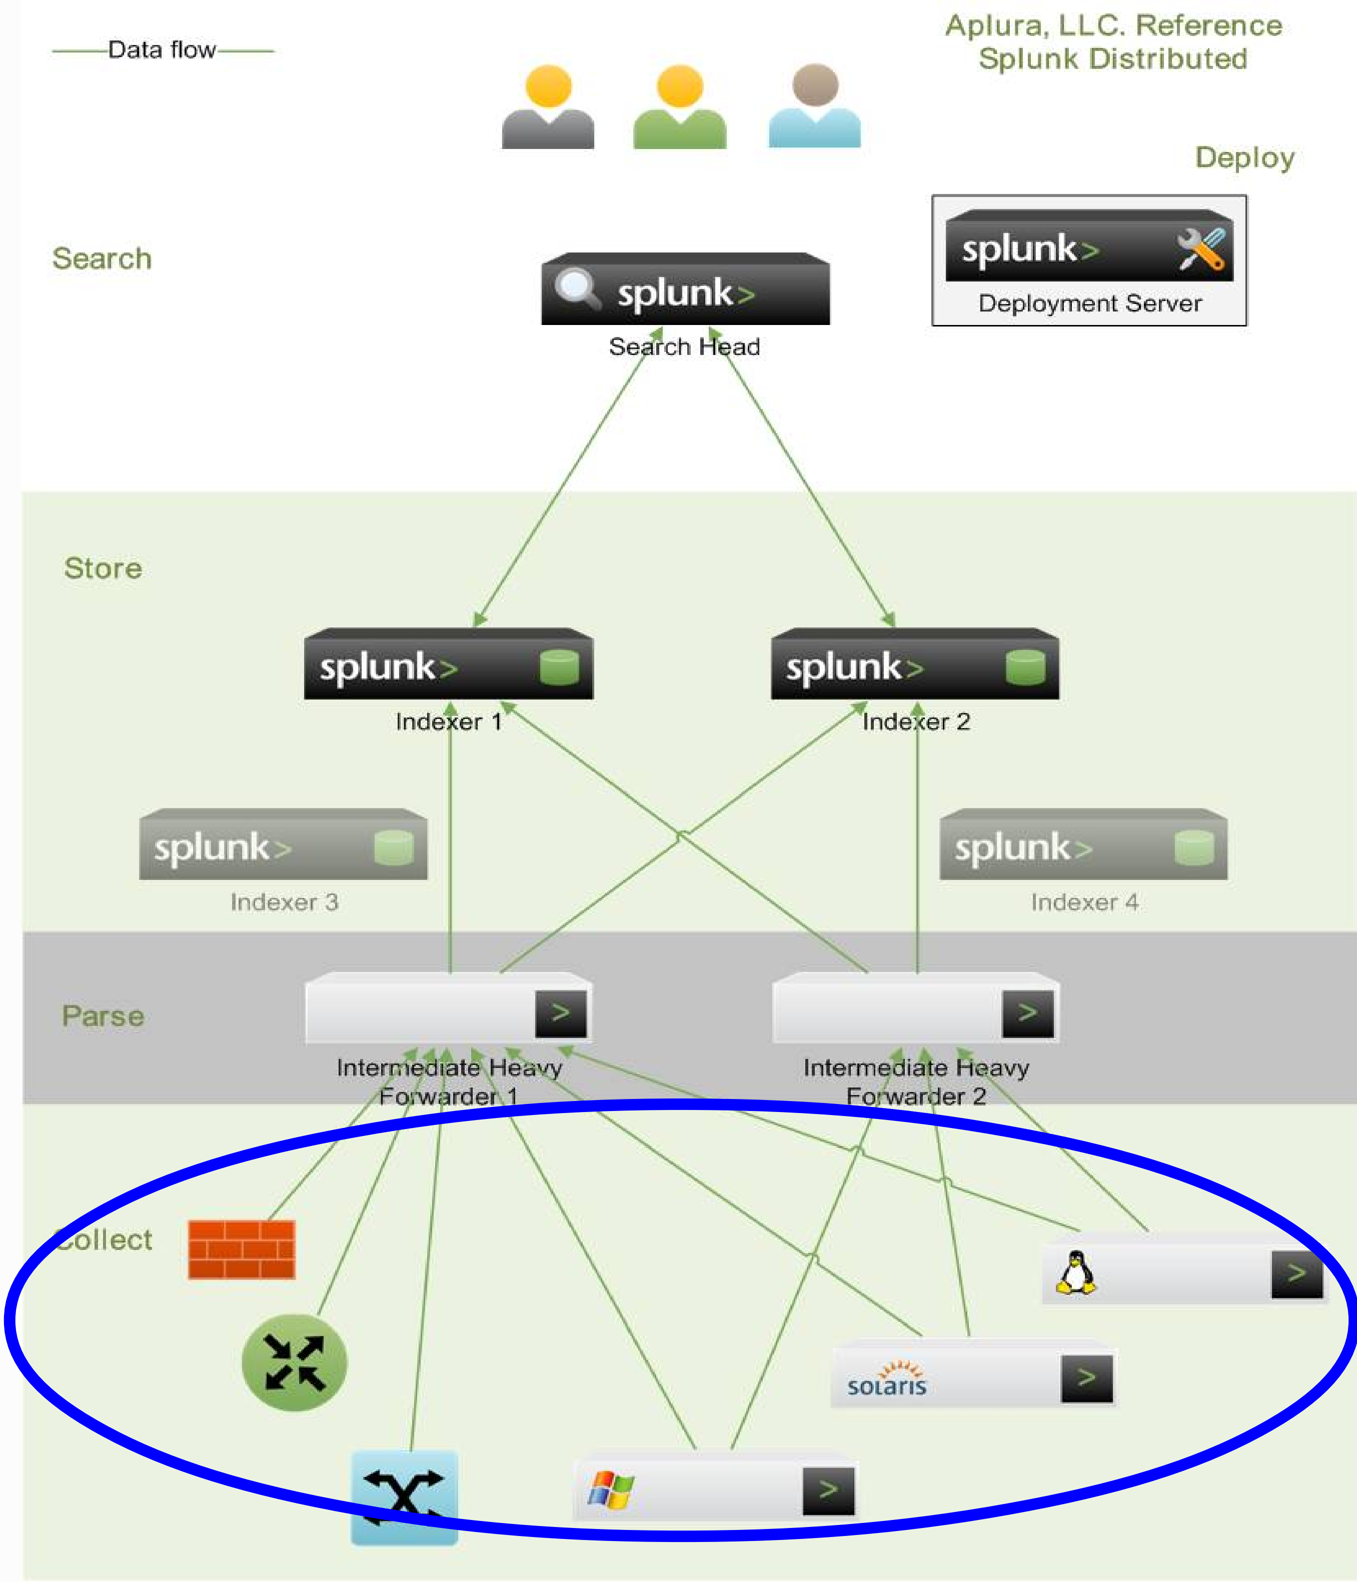

Aplura, LLC. Reference
Splunk Distributed

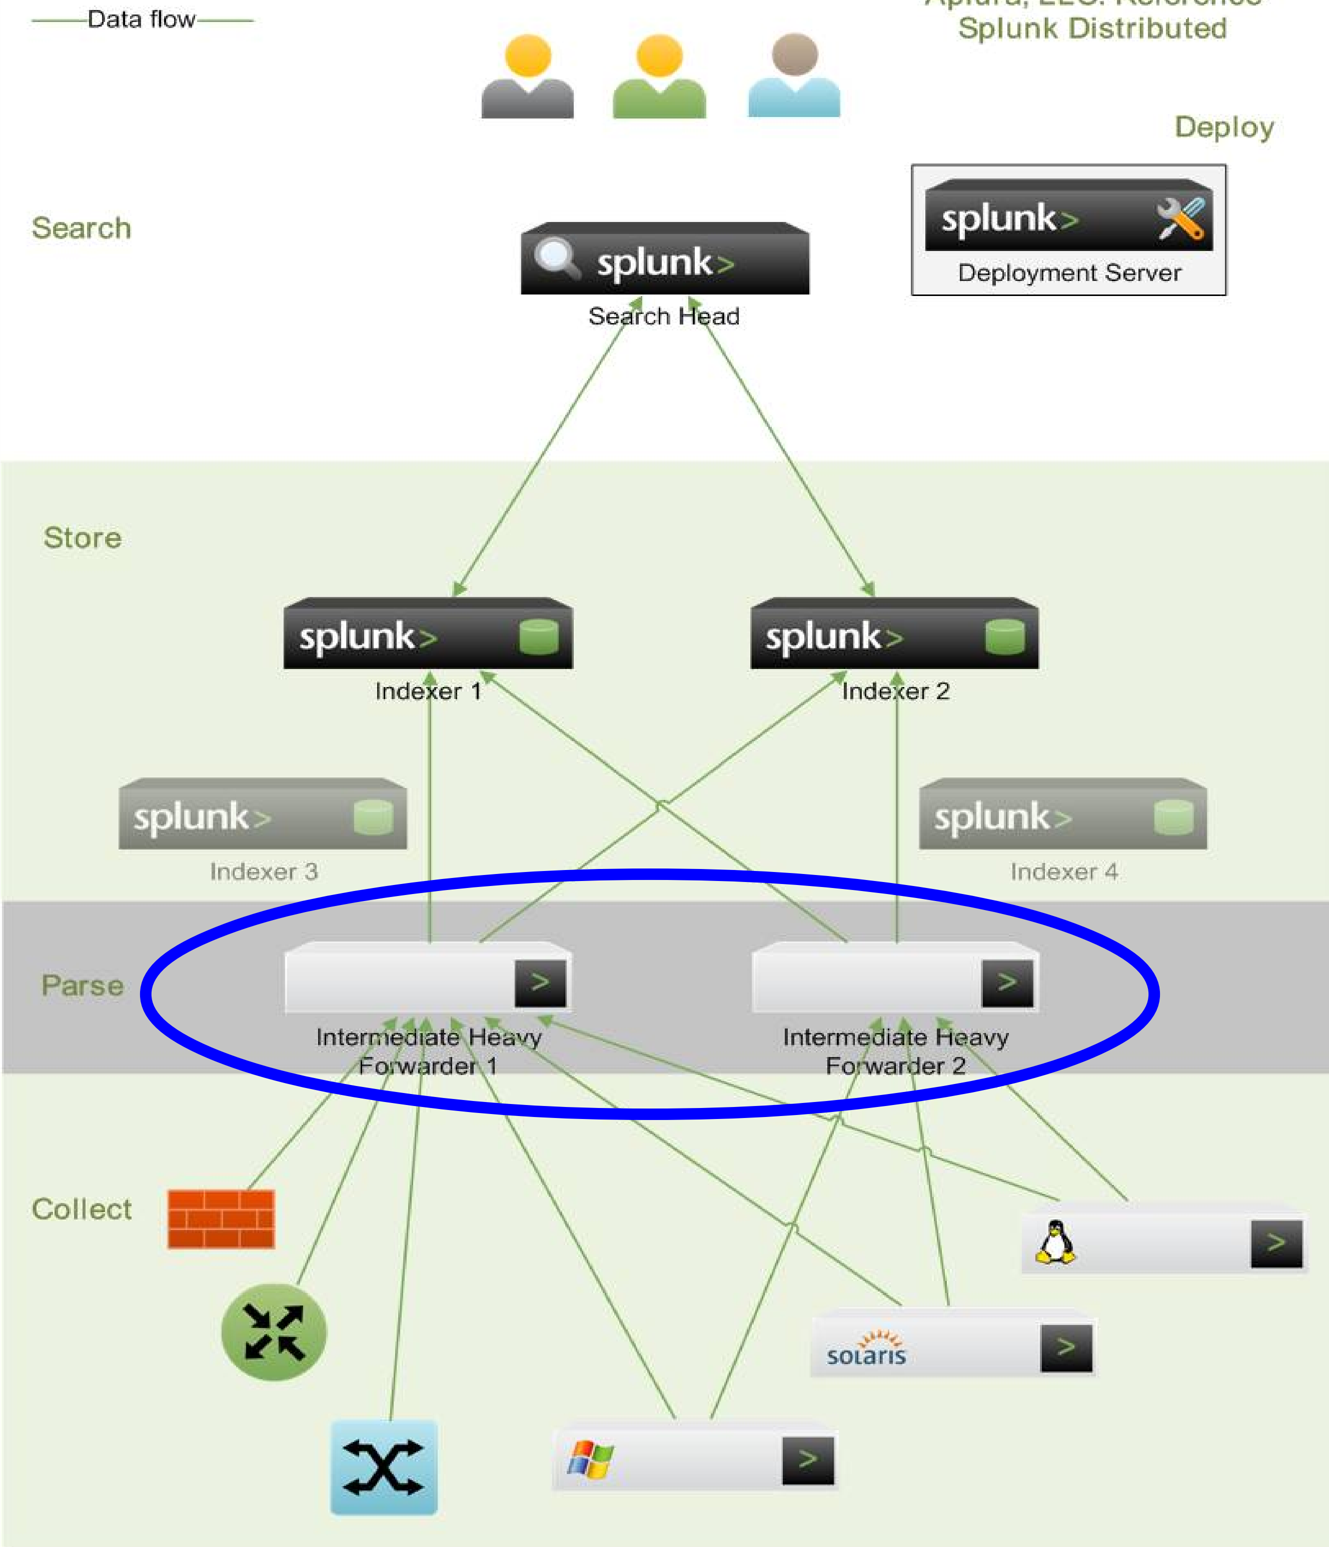

Parsing Operations:

- Sear raw data
- Prep for Index
- Time-stamping
- Event Breaking
- Indexed Fields

IHF:

- Deploy >= 2 for HA
- Low Storage Needs
- VMs Work Great
- Consider Regional or Zone-specific

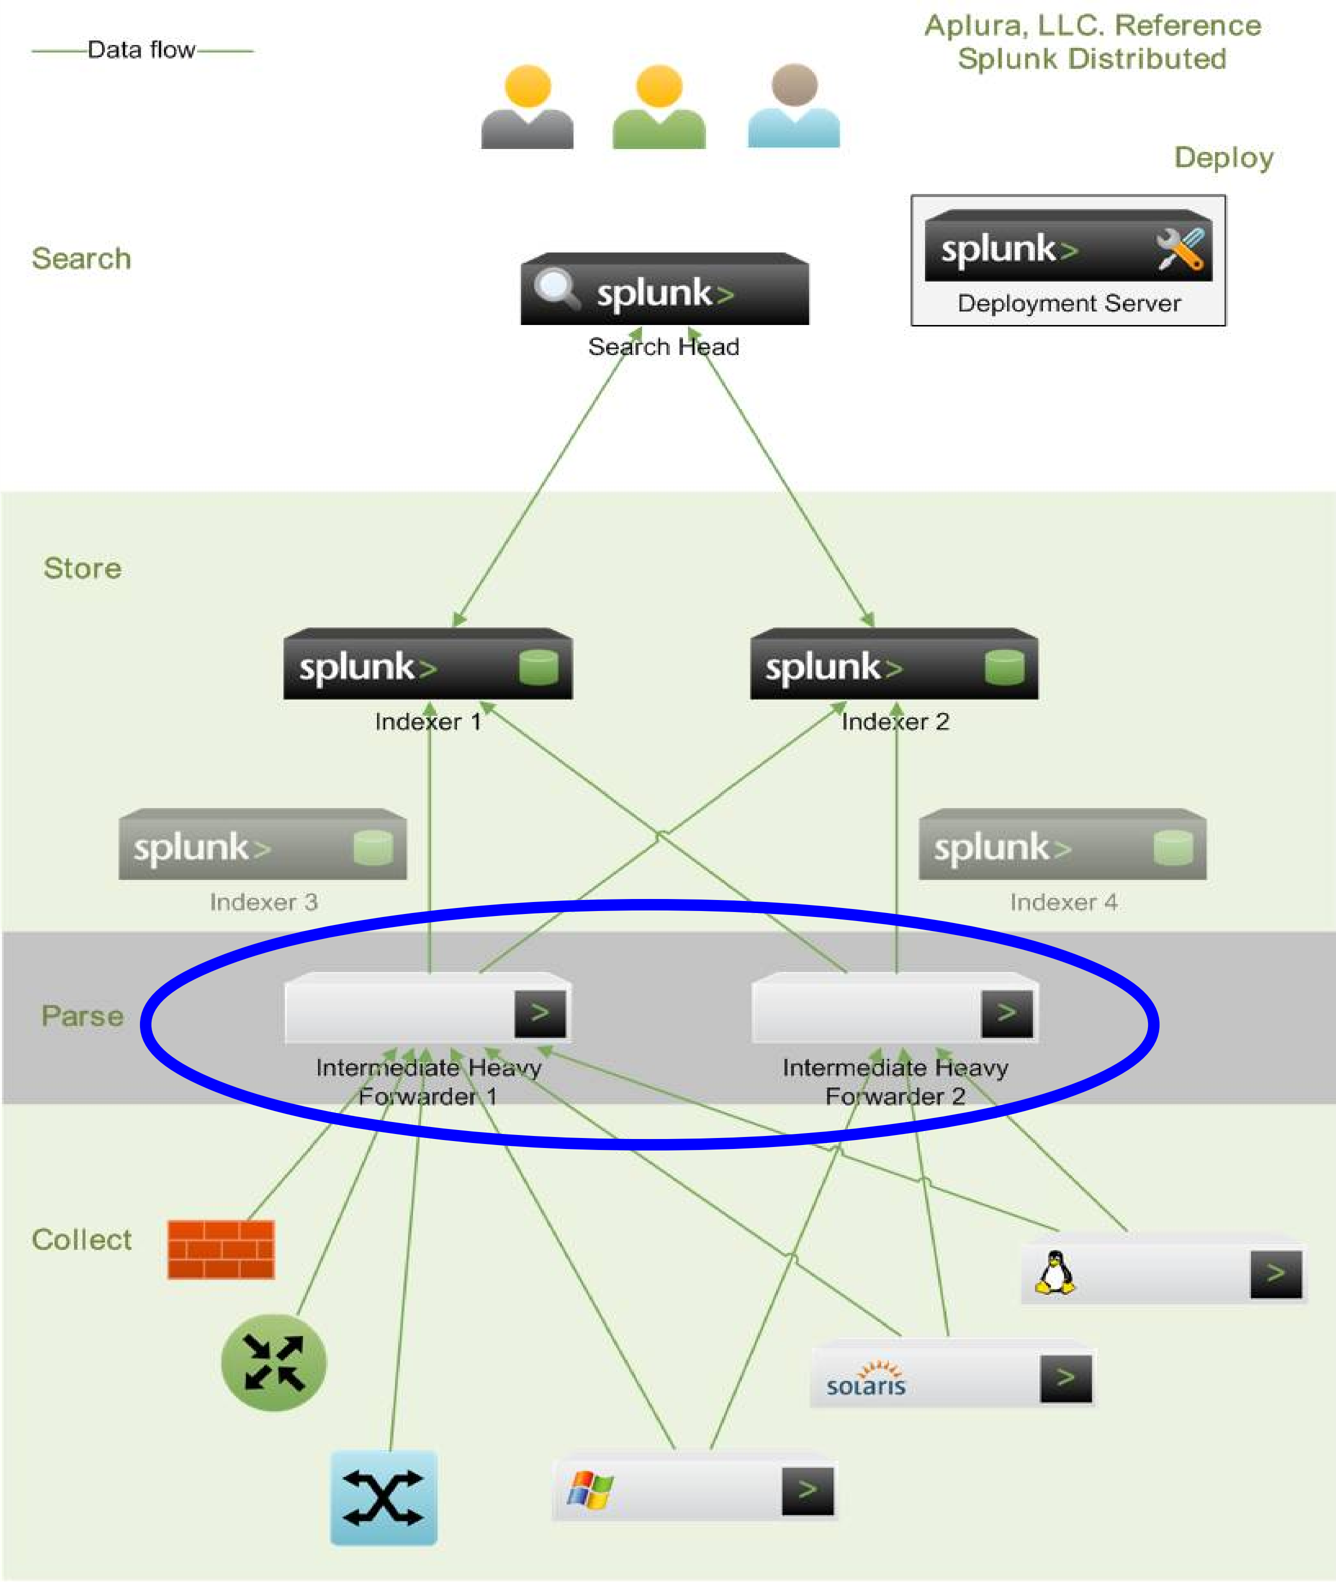

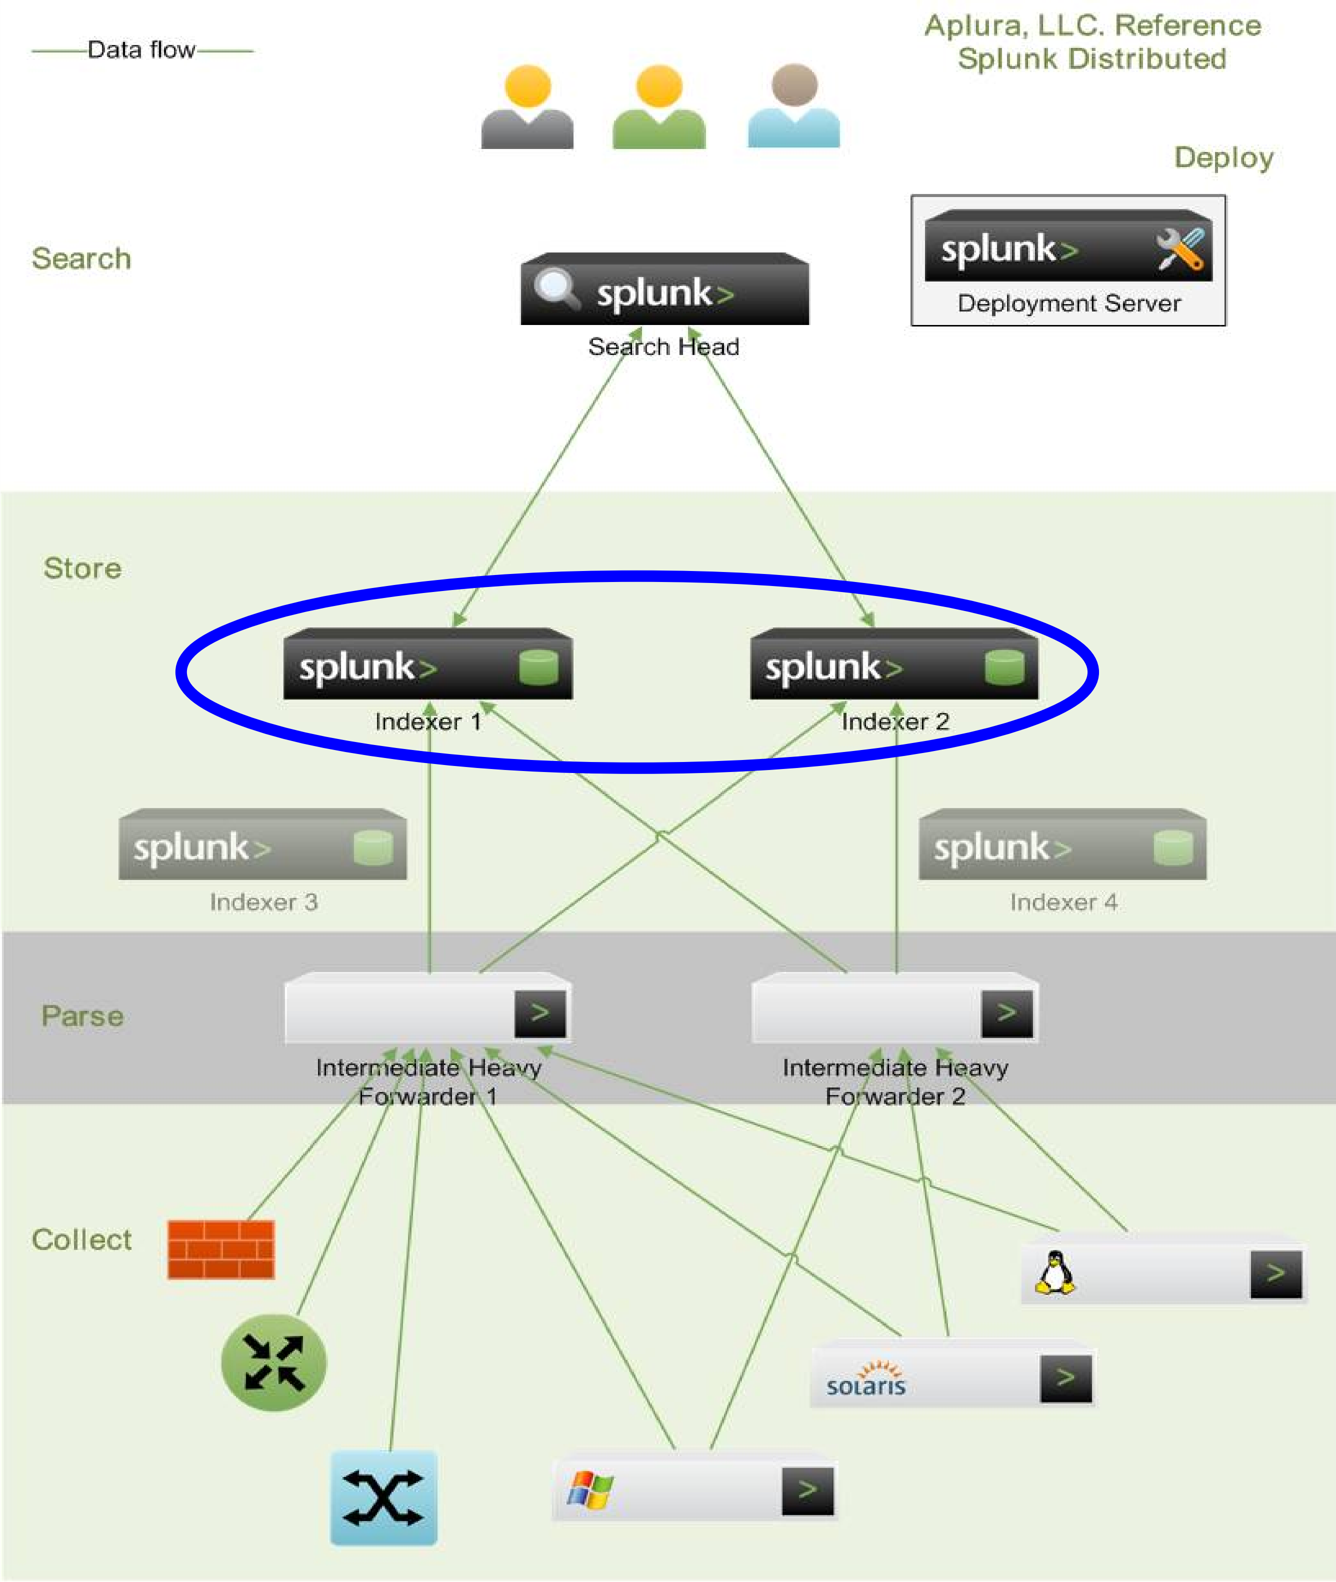

Aplura, LLC. Reference Splunk Distributed

Splunk Indexers:

- IHF Tasks
- Write Data
- Retrieve Data

Plan as a DB

Require 800+ IOPS

VMs are tricky

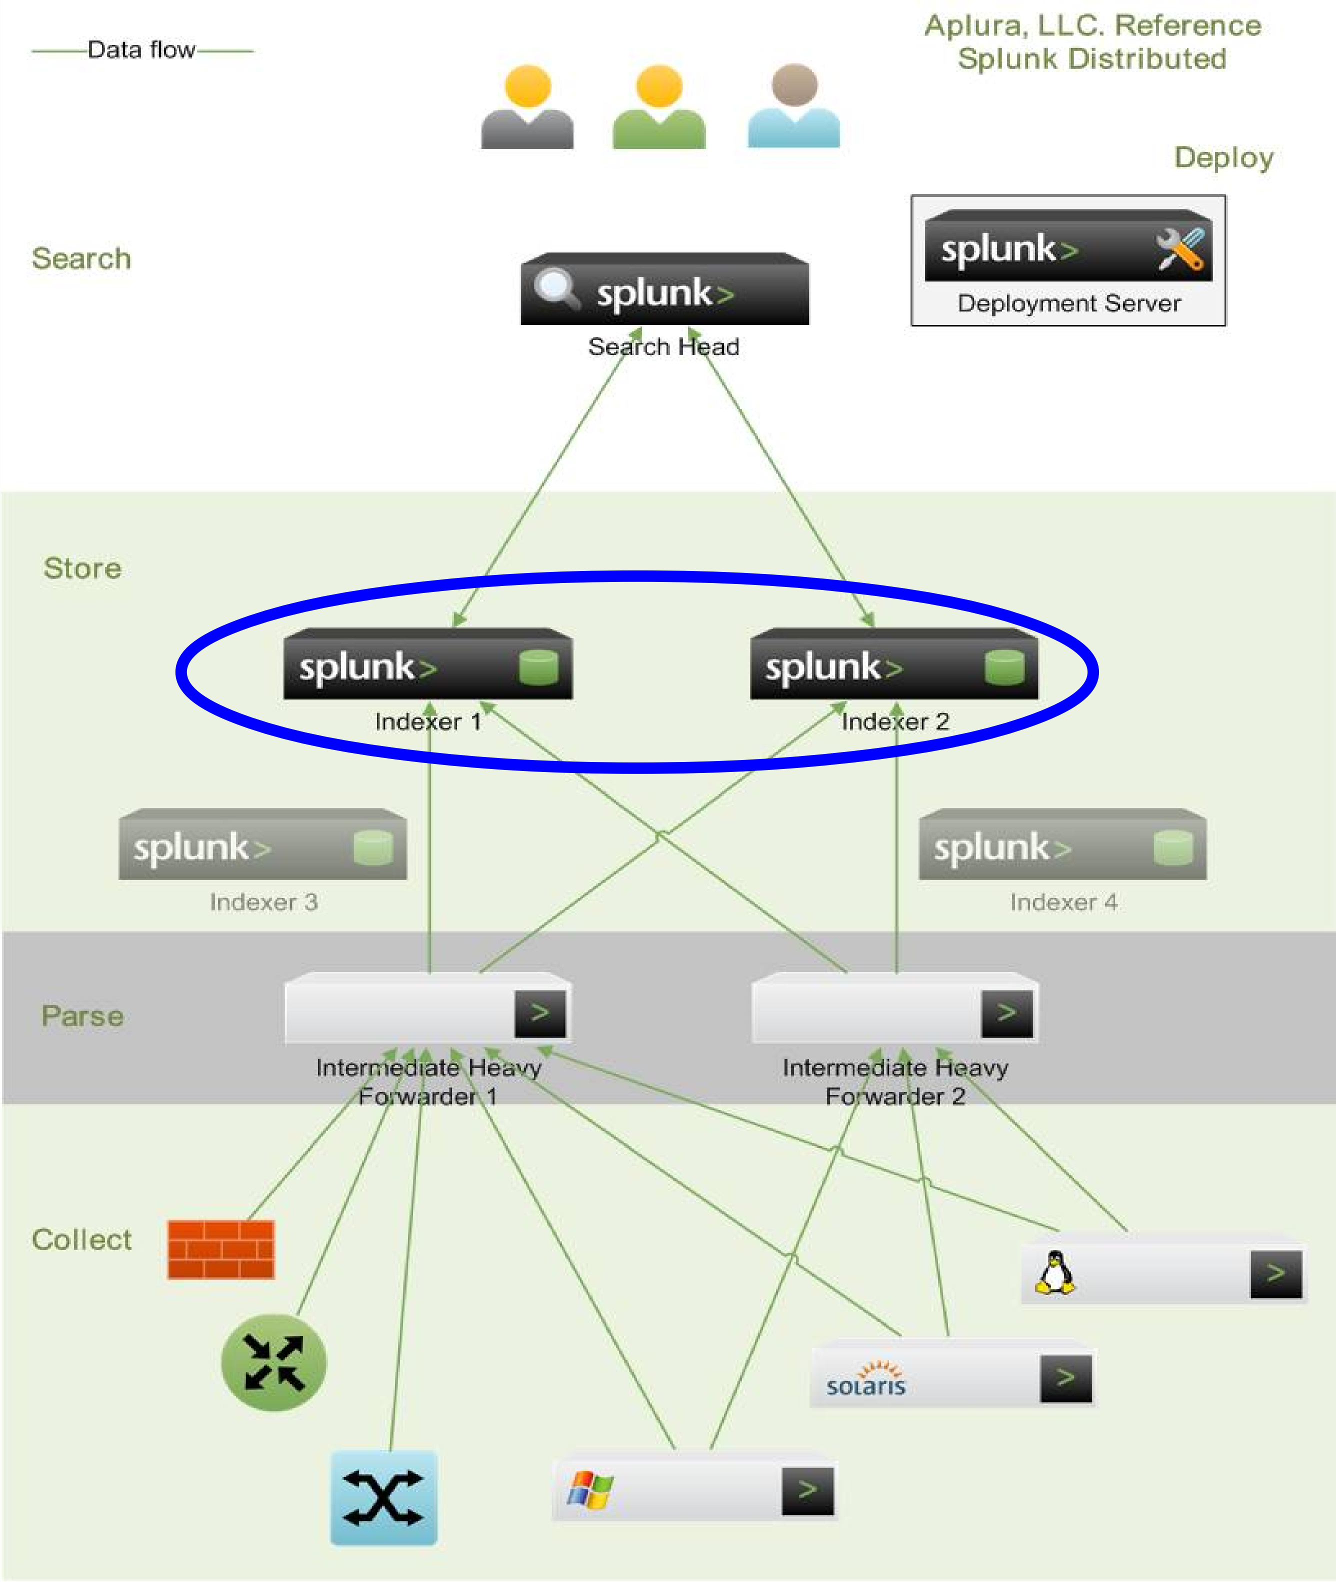

Aplura, LLC. Reference
Splunk Distributed

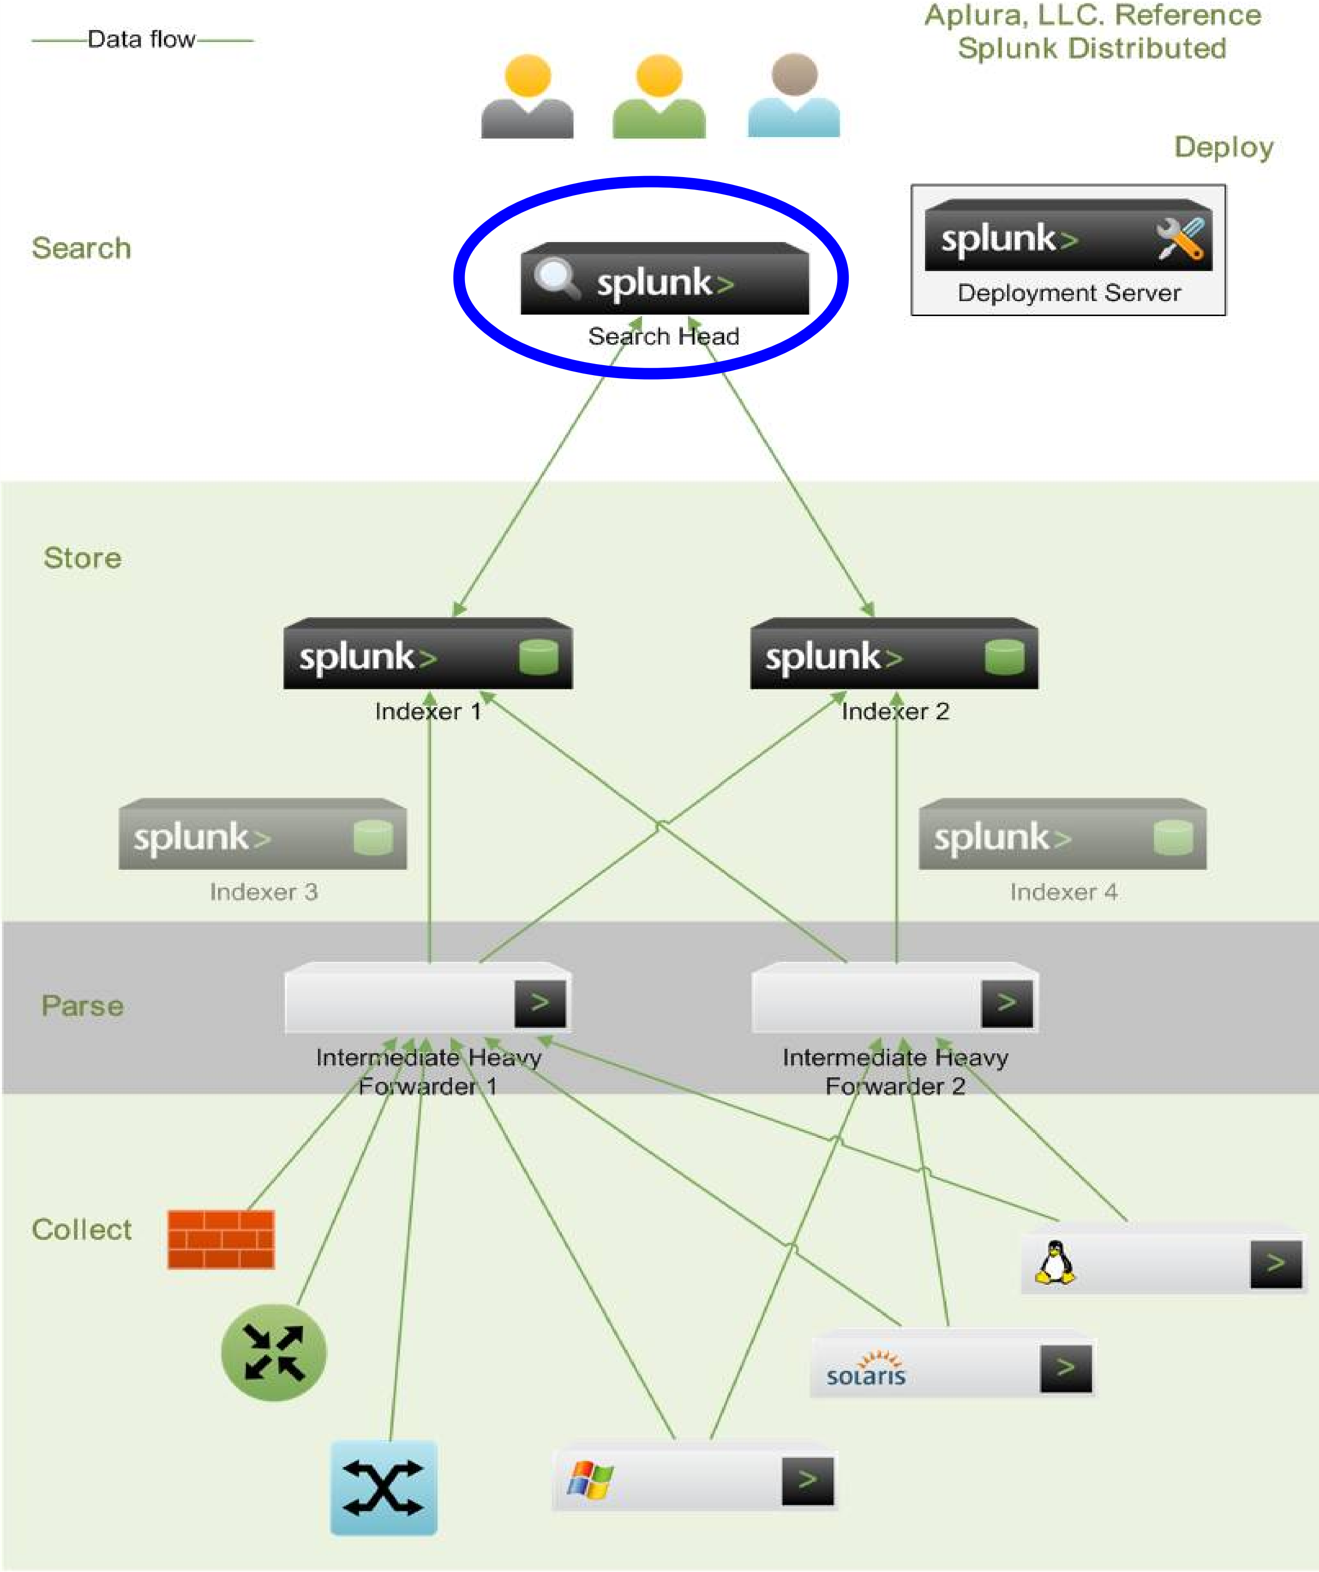

Search Head is where the magic happens

The SH queries the Indexers

Accounts, Permissions, Dashboards, Searches, and Fields, all live in Search

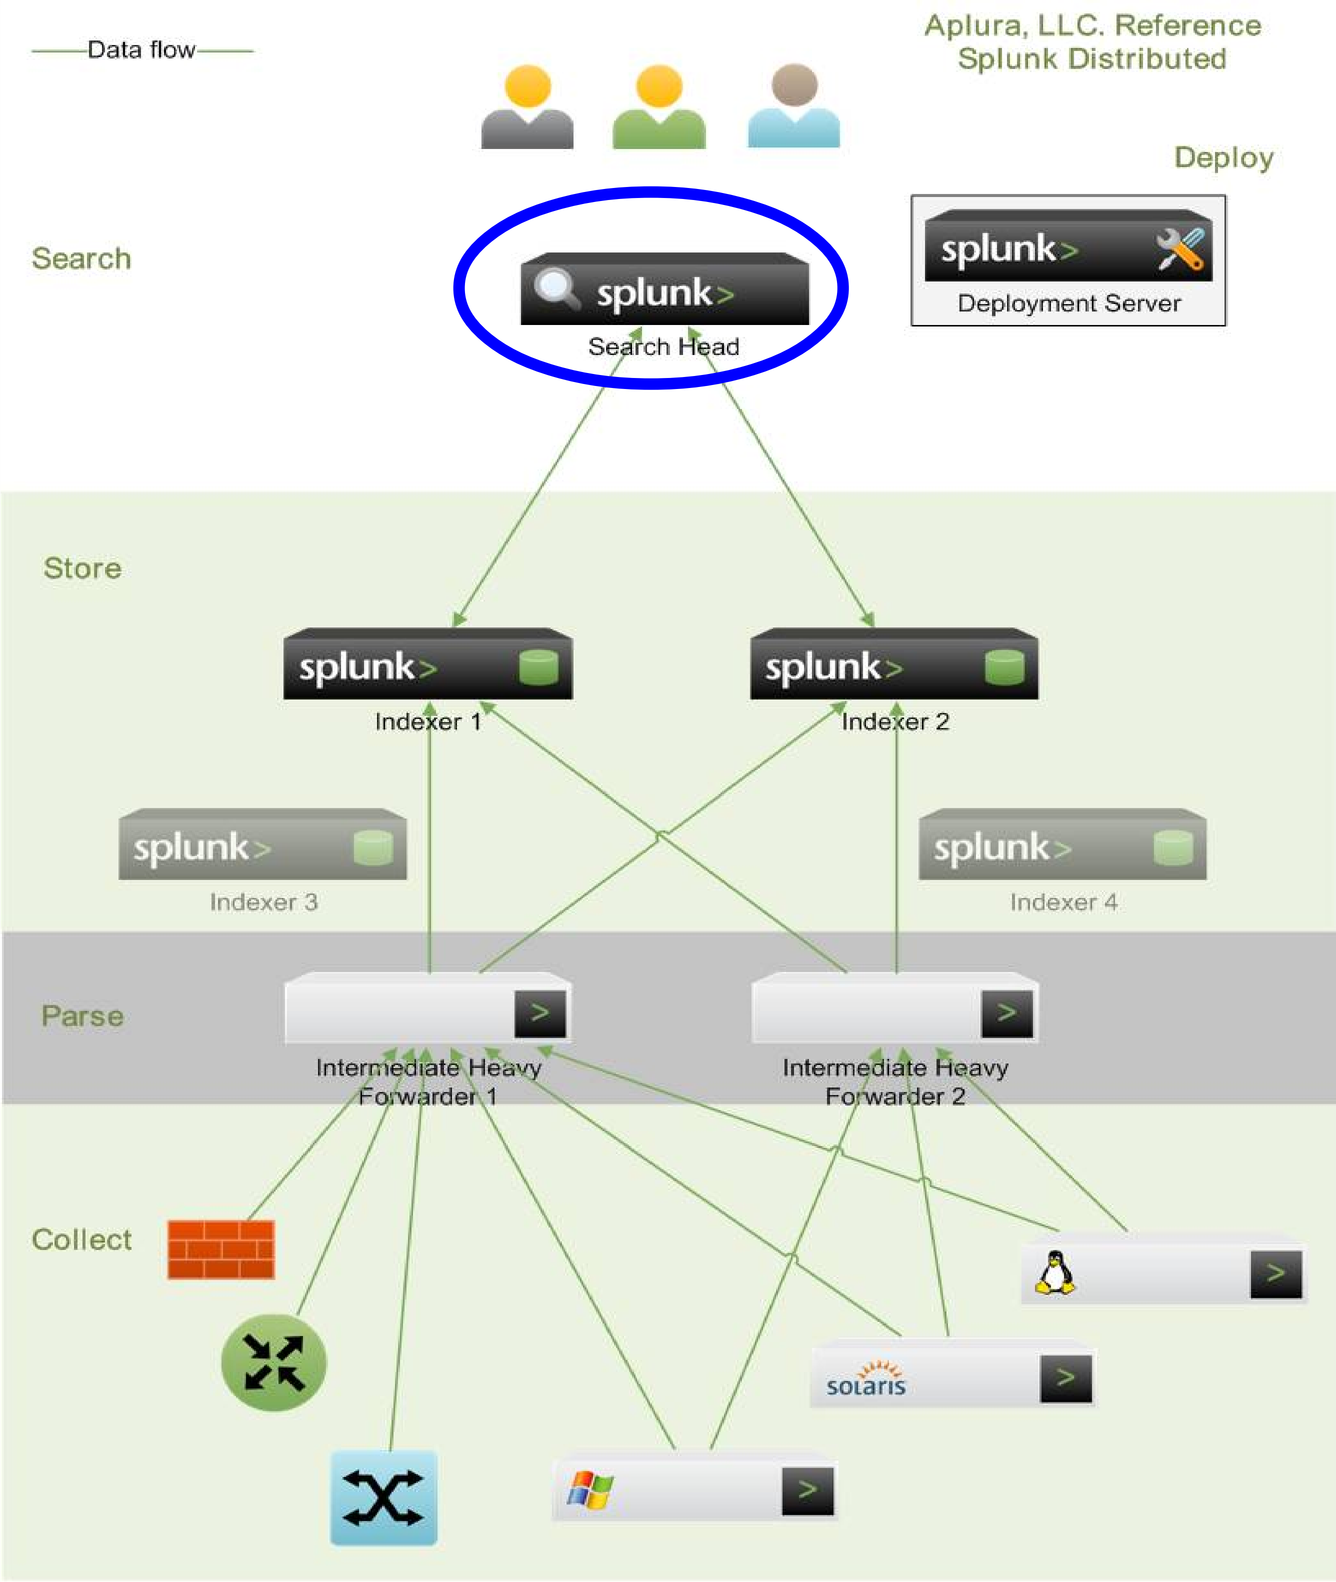

Migrations, Migrations, Migrations!

Migrations, Migrations, Migrations!

This is my favorite part of (\w+)

Warning

- Splunk has many options and is very flexible
- Slides here are to give you ideas
- This may not be best for your use-case
- When in doubt:
    - Read Docs
    - Hit Splunkbase
    - Call Support
    - Call PS

Prep Yourself

- Splunk migrations are easily done
- You can do almost anything without data loss
  - … and minimal downtime
- But, this is your data. This is valuable
- Understand how data flows through Splunk
- Understand the states of data
- Understand how Splunk communicates
- PT: Plan out your steps in advance

Old adage applies well to migrations

You can have it good

You can have it fast

You can have it cheap

Old adage applies well to migrations

You can have it good

You can have it fast

You can have it cheap

Choose any two

Migrate Solo to Distributed

Migrate Solo to Distributed

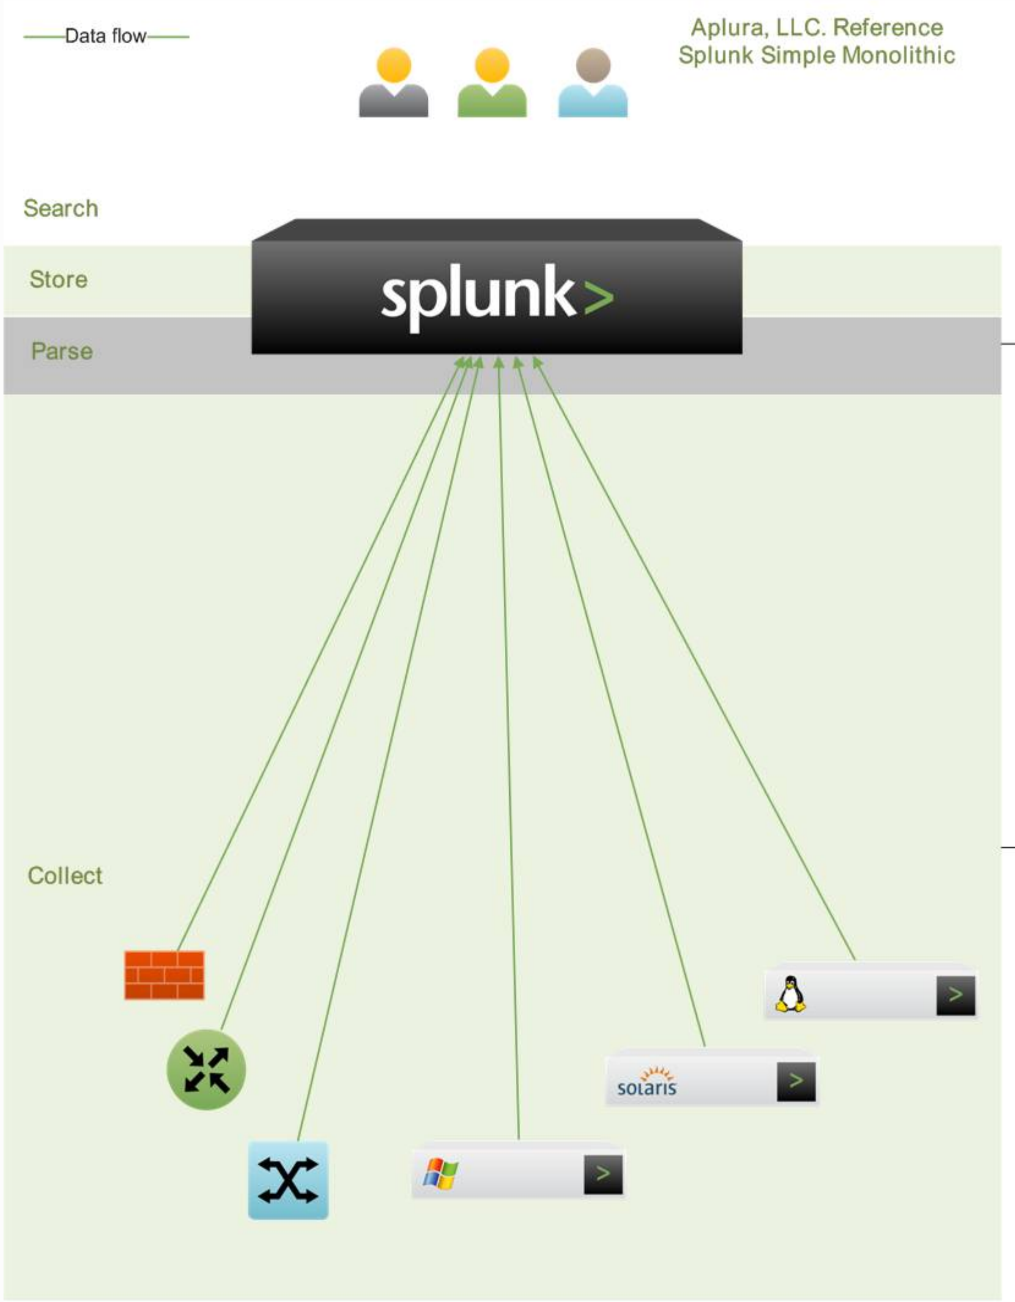

Migrate Solo to Distributed

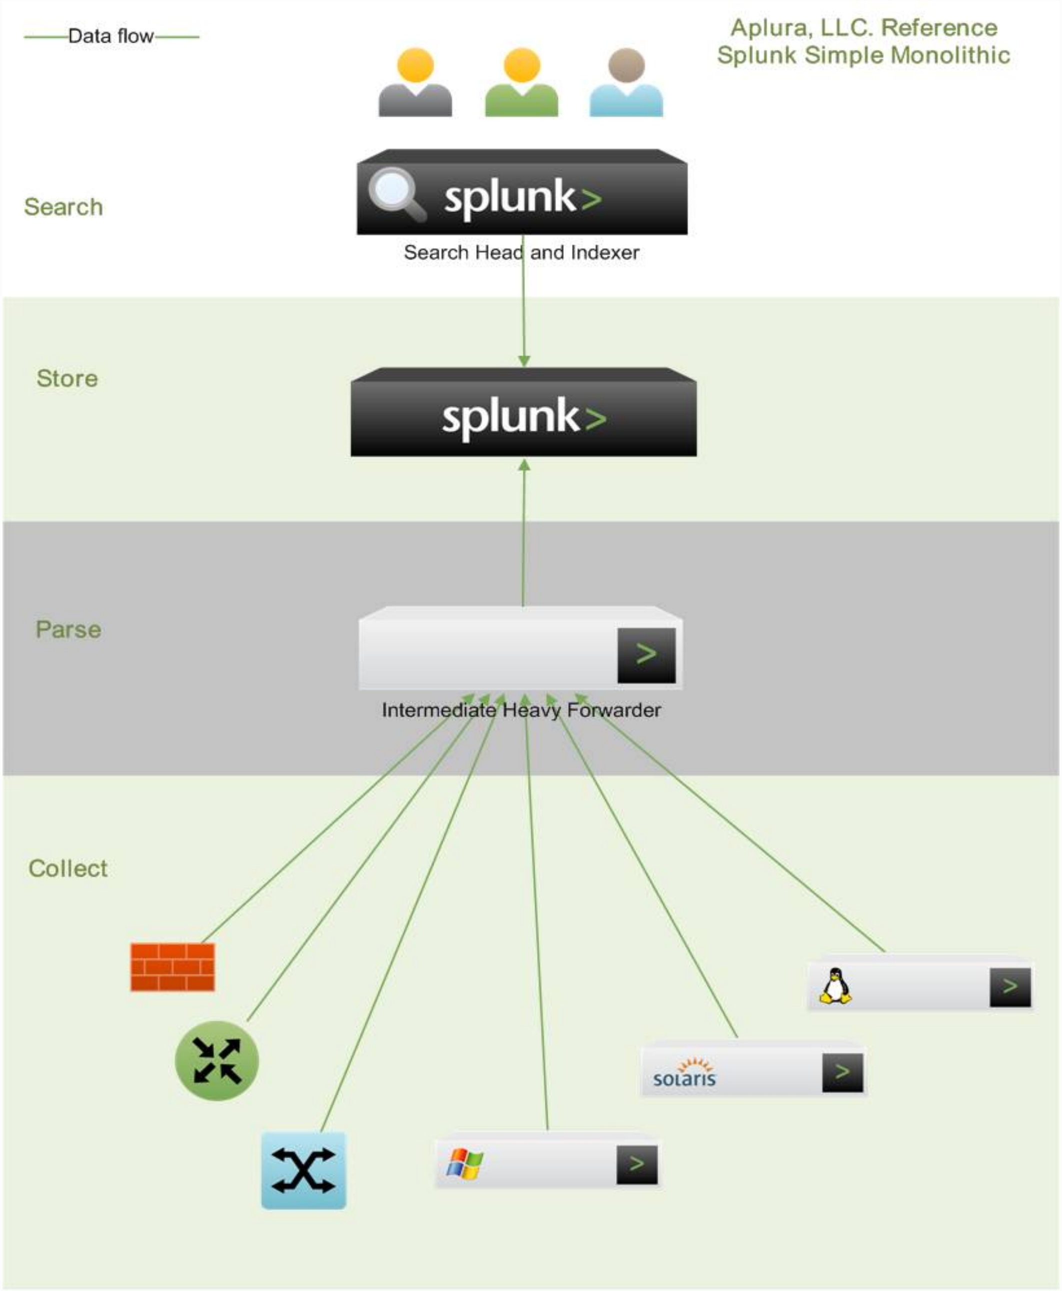

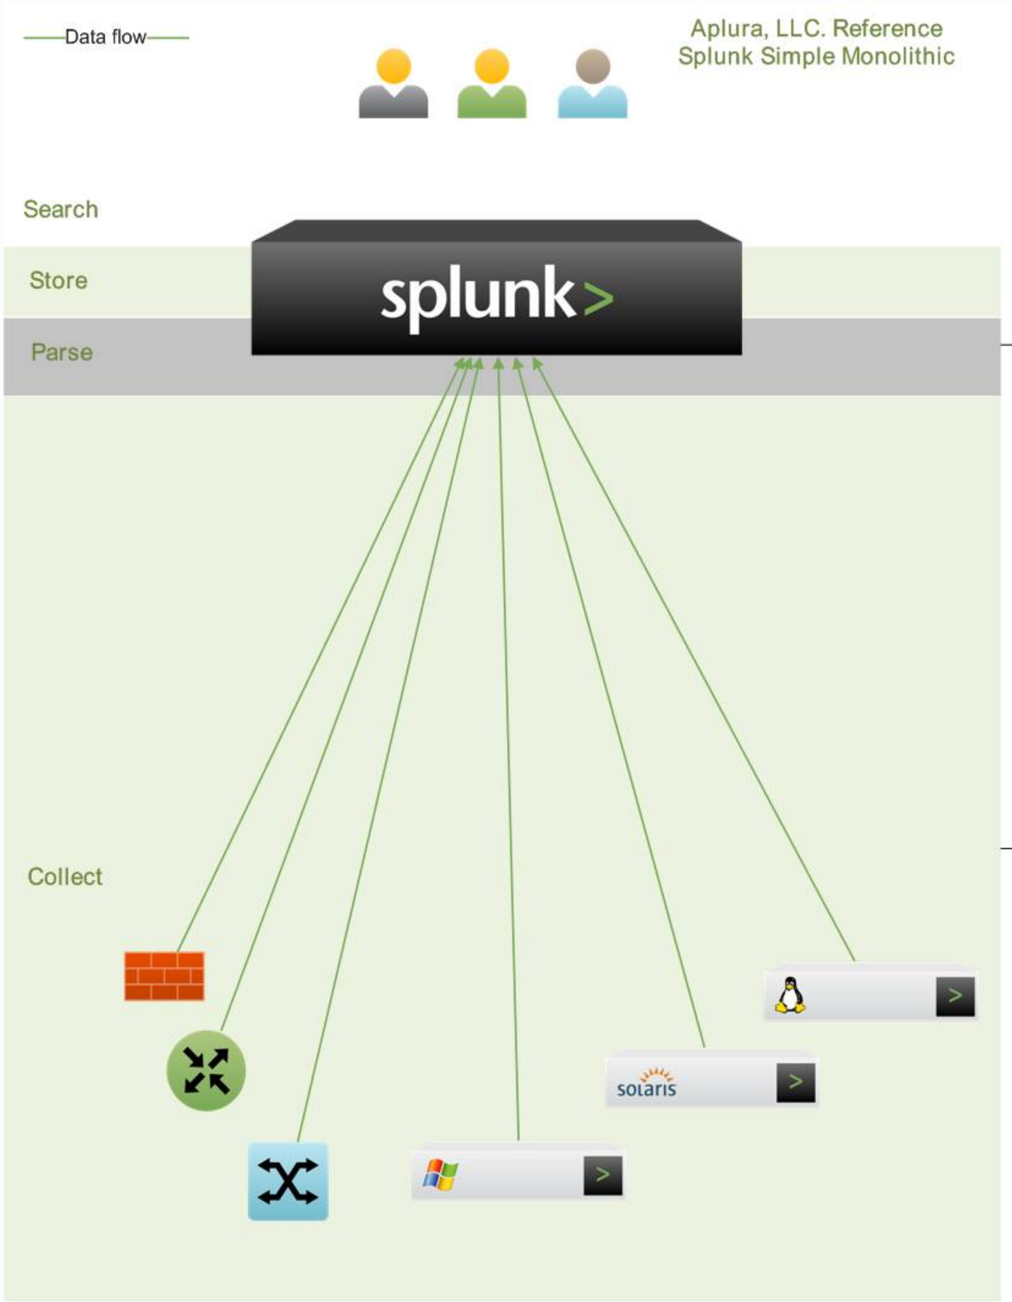

Migrate Solo to Distributed

- No data was harmed
- Steps:
  - PT: Consider Splunk DS
  - PT: Separate Parsing
  - Move Search Tier
  - Add another indexer
    - Rinse, repeat.

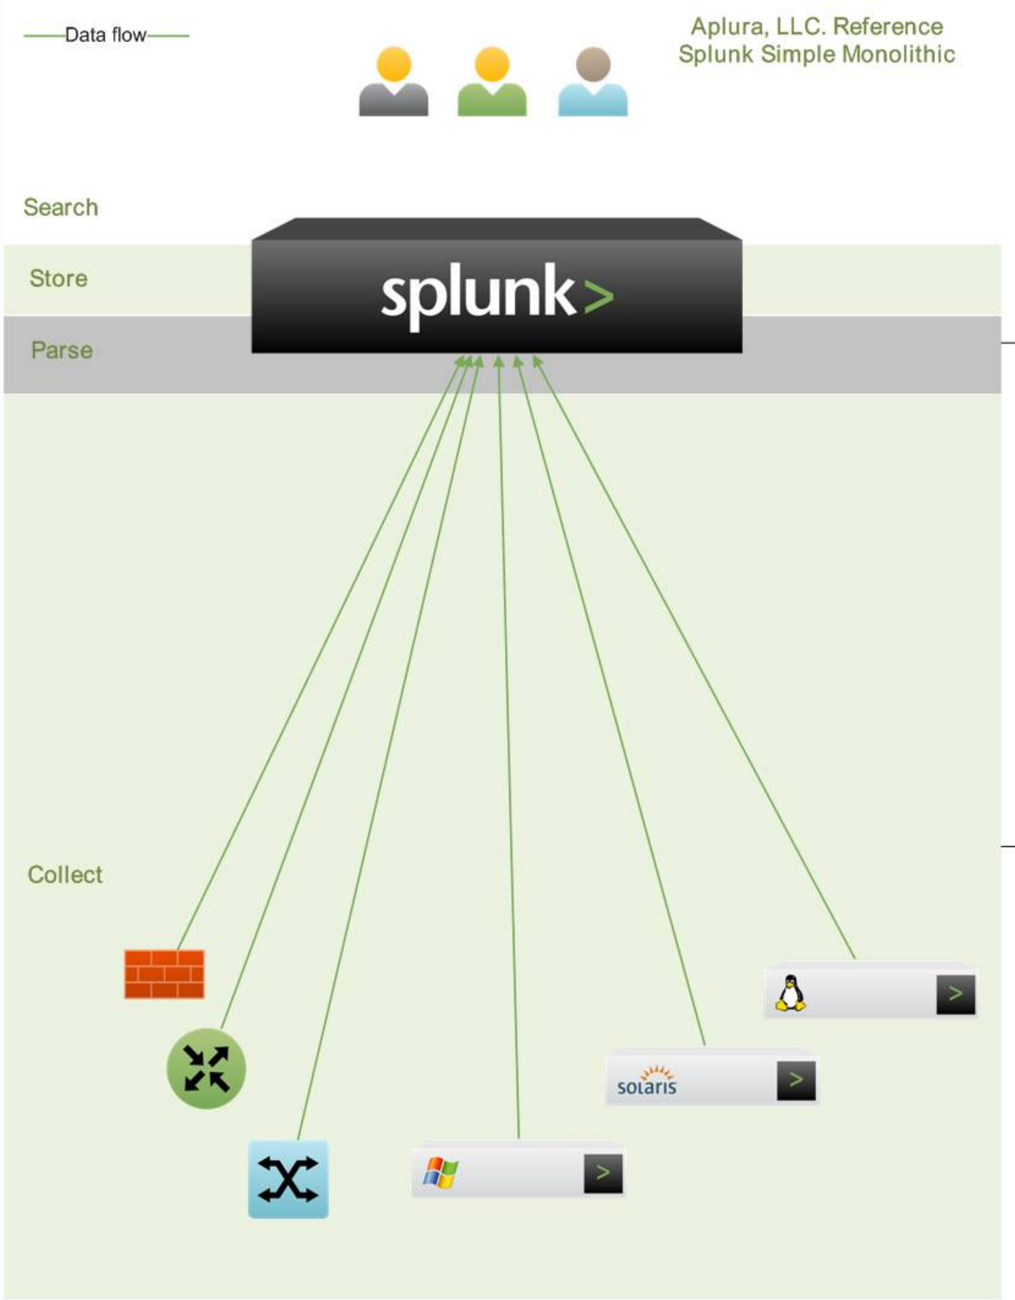

Migrate S2D: Separate Parsing

- Controlling the parsing is an important first step
- Consider an Intermediate Heavy Forwarder
  - Allows control of the data-flow
  - Discrete apps to manage parse-time rules
  - Brokers TCP sessions to Universal Forwarders
  - Reduces load on indexers
  - It is like hitting the “turbo button”
- Often done with no downtime and no data loss

Migrate S2D: Separate Parsing

- Setup Splunk on new IHF host
  - Sparse config, no index, forwards all data to Solo1
- Sync all parse-time configs to IHF from Solo1
- Have UFs send data to new IHF
  - If possible, update senders with new receiver
- Now UF-->IHF-->Solo1

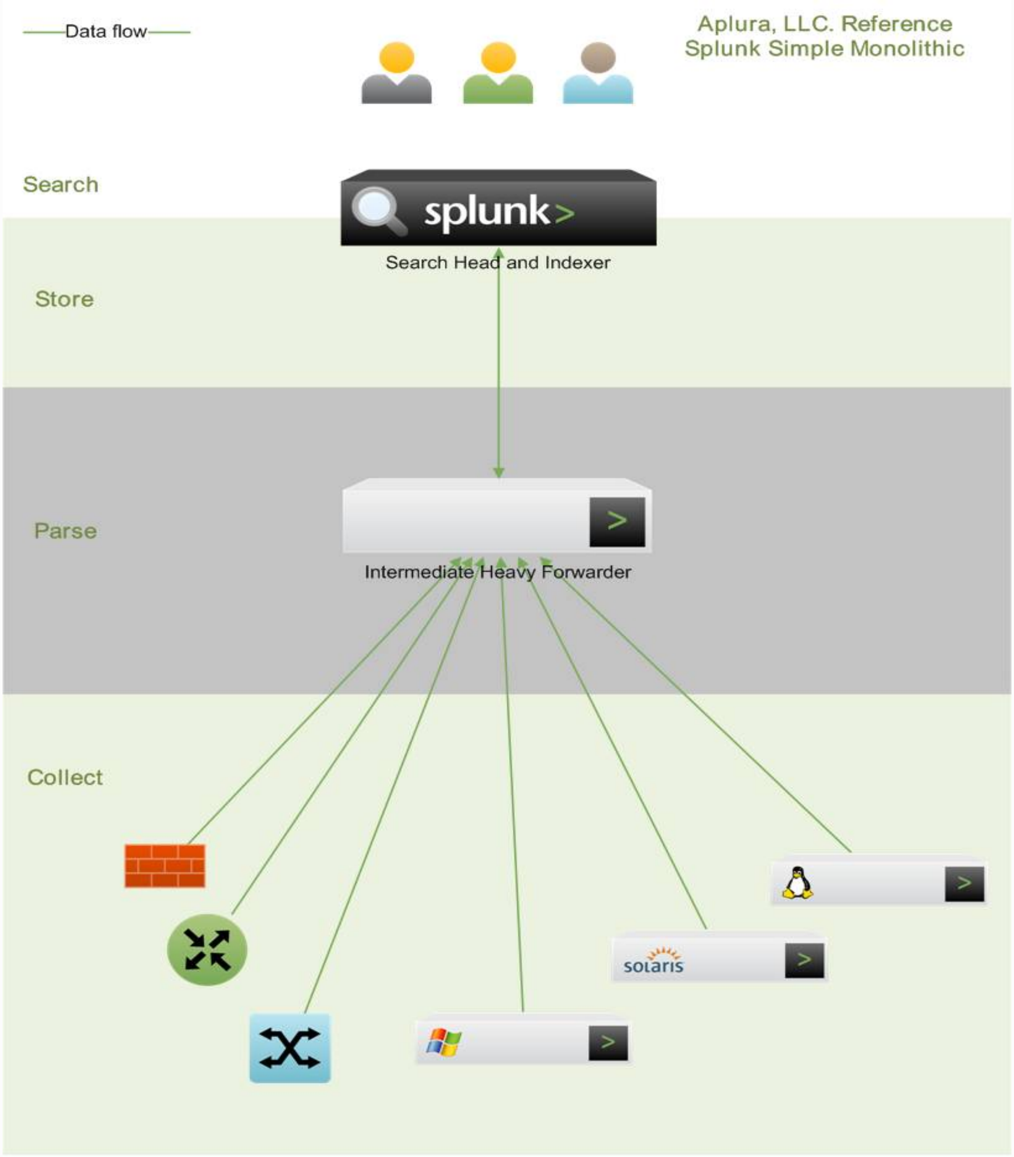

Migrate S2D: Dist Search

- Separate Search, allows for horizontal scaling
- This affects the user environment
- Should plan for light user downtime

Migrate S2D: Dist Search

- Setup Splunk on Search
  - No local indexing
- Sync all search configs from Solo1 to SH1
  - Don't forget User configs
  - Use opportunity to simplify with discrete apps
  - PT: Watch for app name precedence
- Make DNS change of CNAME used for Splunk
- PT: Disable/remove searches from Solo1

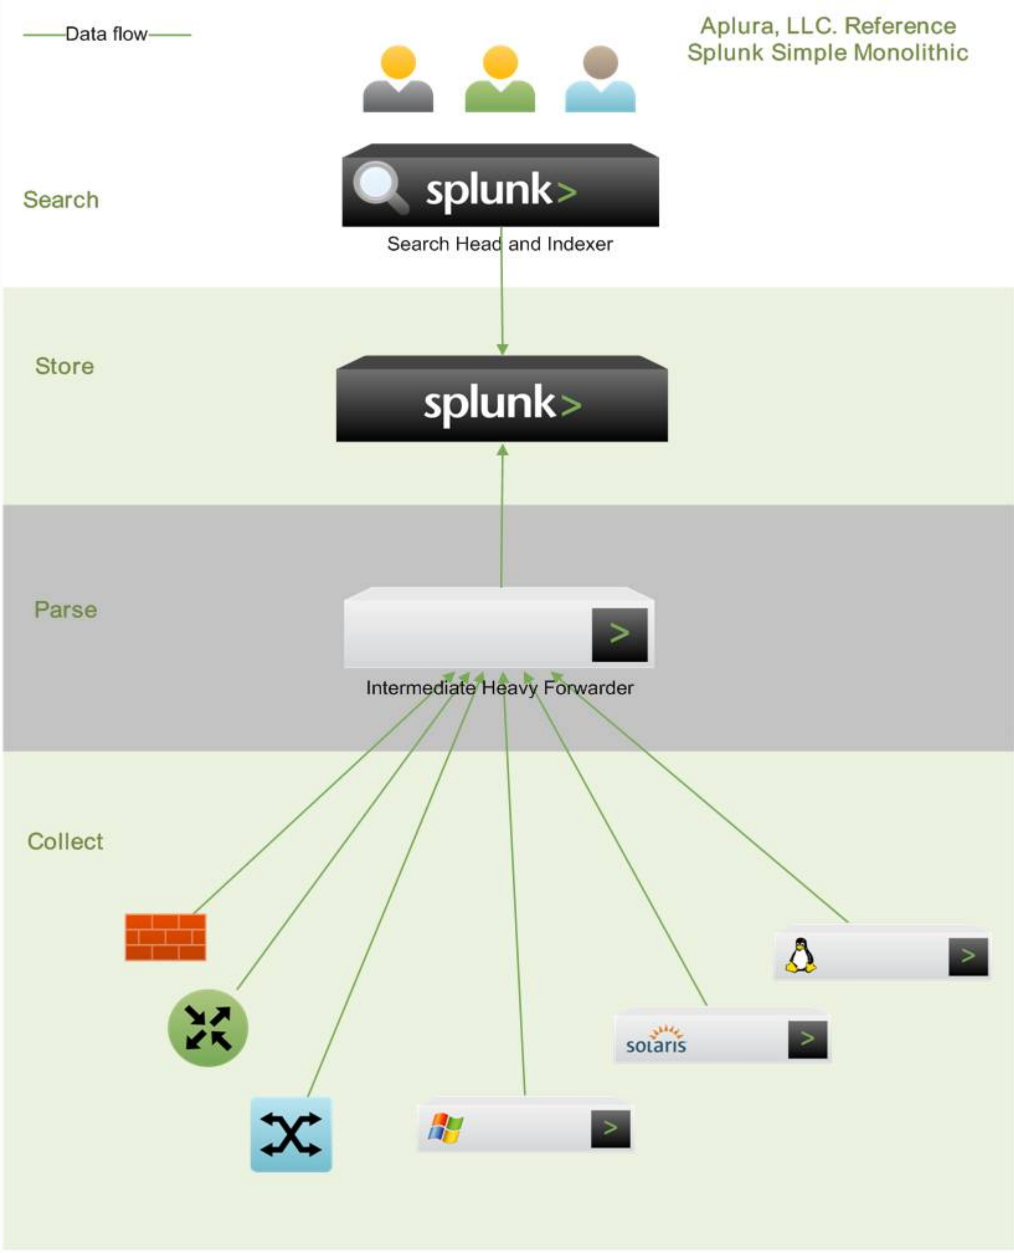

Migrate Indexer

- Indexers can be moved, with proper planning
- You must consider some important things
  - When on, they are like databases (so, be careful)
  - When off, it is all flat-files (easily moved or rsync)
- Depending on specific objective, consider options (what is Splunk capable of)

Migrate Indexer - Steps

- Build new Indexer
- Ready to Migrate:
  - Sync configs from Old1 to New1
  - Stop data from going to Old1
  - Old1: Restart Splunk, then stop (rolls buckets)
  - Rsync hot/warm, cold buckets from Old1 to New1
  - Restart Splunk on New1
  - Add New1 to SH (test to see _internal on New1)
  - Redirect IHF to send to New1

Merge Index

- Splunk data is indexed in “buckets”
- Together, the buckets make up an index
- What if you want to consolidate indexes?

Merge Index – Bucket Brigade

- Hot buckets are live (currently being written to)
  - Never touch these...ever!
- Warm buckets are read-only
  - PT: Rsync in-advance of migration date to prep some warm buckets and shorten outage window
- The trick with buckets is not to confuse Splunk
  - Rename the bucket IDs appropriately

Merge Index – Bucket Brigade

- Examine a bucket: $SPLUNK_HOME/var/lib/splunk/INDEX/db

```
-rw------- 1 root root  309 May 14 22:59 .bucketManifest
-rw------- 1 root root   10 May 14 22:34 CreationTime
drwx--x--x 3 root root 4096 May 14 22:53 db_1337049851_1337048652_0
drwx--x--x 3 root root 4096 May 14 22:57 db_1337050398_1337049799_1
drwx--x--x 3 root root 4096 May 14 22:58 db_1337050617_1337050018_2
drwx--x--x 2 root root 4096 May 14 22:34 GlobalMetaData
-rw------- 1 root root  101 May 14 22:59 Hosts.data
drwx--x--x 3 root root 4096 May 14 23:00 hot_v1_3
-rw------- 1 root root  302 May 14 22:58 .metaManifest
-rw------- 1 root root  276 May 14 22:59 Sources.data
-rw------- 1 root root  109 May 14 22:59 SourceTypes.data
```

Merge Index – Bucket Brigade

- Steps to merge indexes
  - Backup everything first
  - Stop data-flow into Indexer(s)
  - Restart Splunk on Indexer(s) (forces bucket roll)
  - Sync buckets from old into new
  - Rename buckets such that bucketIDs increment
    - PT: Use Larry Wall's rename.pl
  - Purge ".bucketmanifest" forcing splunk to rebuild
  - Restart Splunk on Indexer new
  - Disable Index on old
  - Start data-flow

Merge Index - Rename Script

- I use Larry Wall's "rename" perl script (comes with Ubuntu but NOT CentOS/RHEL)
- Larry Wall's "rename" takes sed-style matches, so the following will work:

```
/bin/sh
cd $SPLUNK_DB/TARGETINDEX/(db|colddb)
START=000
for i in ls -rtd db_*; do START=$(($START + 1)); rename -nv "s/d+$/$START/" $i ;done
```

- The ls with rtd args, sorts reverse by time and only shows the directories, so the oldest bucket will be listed first and therefore will be the first bucket in your list to keep the order somewhat the way it would have been originally.
- The business with the START thing helps to maintain a counter so you can move through them one at a time. This is important when you are doing both the HOT then COLD DBs, and therefore you need to start on a specific number.

License Master Migration

- License master was added with 4.2
- It uses a splunkd service (tcp/8089)
- In rare cases you might want to move the LM
- Splunk doesn't like when the LM moves
  - Your systems may flip-out temporarily
- Steps vary based on source and dest of LM
- Having a Deployment Server helps a great deal

High-Capacity Collection

- Direction depends on use-case
- Scripted Input
  - Good for light output, not heavy/fast data
  - If heavy/fast, have same script log to file then Splunk monitor that file
- Syslog
  - Ok for light output, but consider a syslog service
  - A syslog service is lighter weight with fewer configs and updates
  - Syslog writes to files, then Splunk monitors. Built-in reliability.
  - See BP guide about kernel/service tuning. This makes a big difference.

HA/DR

- Splunk does HA very well
  - Can be built in to every element of the Architecture
- DR is not intuitive
  - No event replication
- Define what HA/DR mean and mean to you
  - Ensure your terms match the features

HA Architecture

- Forwarders natively can load-balance outputs
- Search Heads can be *pooled* together
- A single outage of Indexer1
  - Results in that stored data not available for queries
  - No loss of new data since incoming will go to Indexer2

DR Architecture

- No built-in and *index-aware* event replication
- You have several options
  - Have Splunk forward data to alternate site
  - Use OS tools to sync data to alternate site
  - Use Storage to replicate data to alternate site
- None of these are seamless
- This topic is an entire presentation in itself

DR Architecture – Splunk Forwards

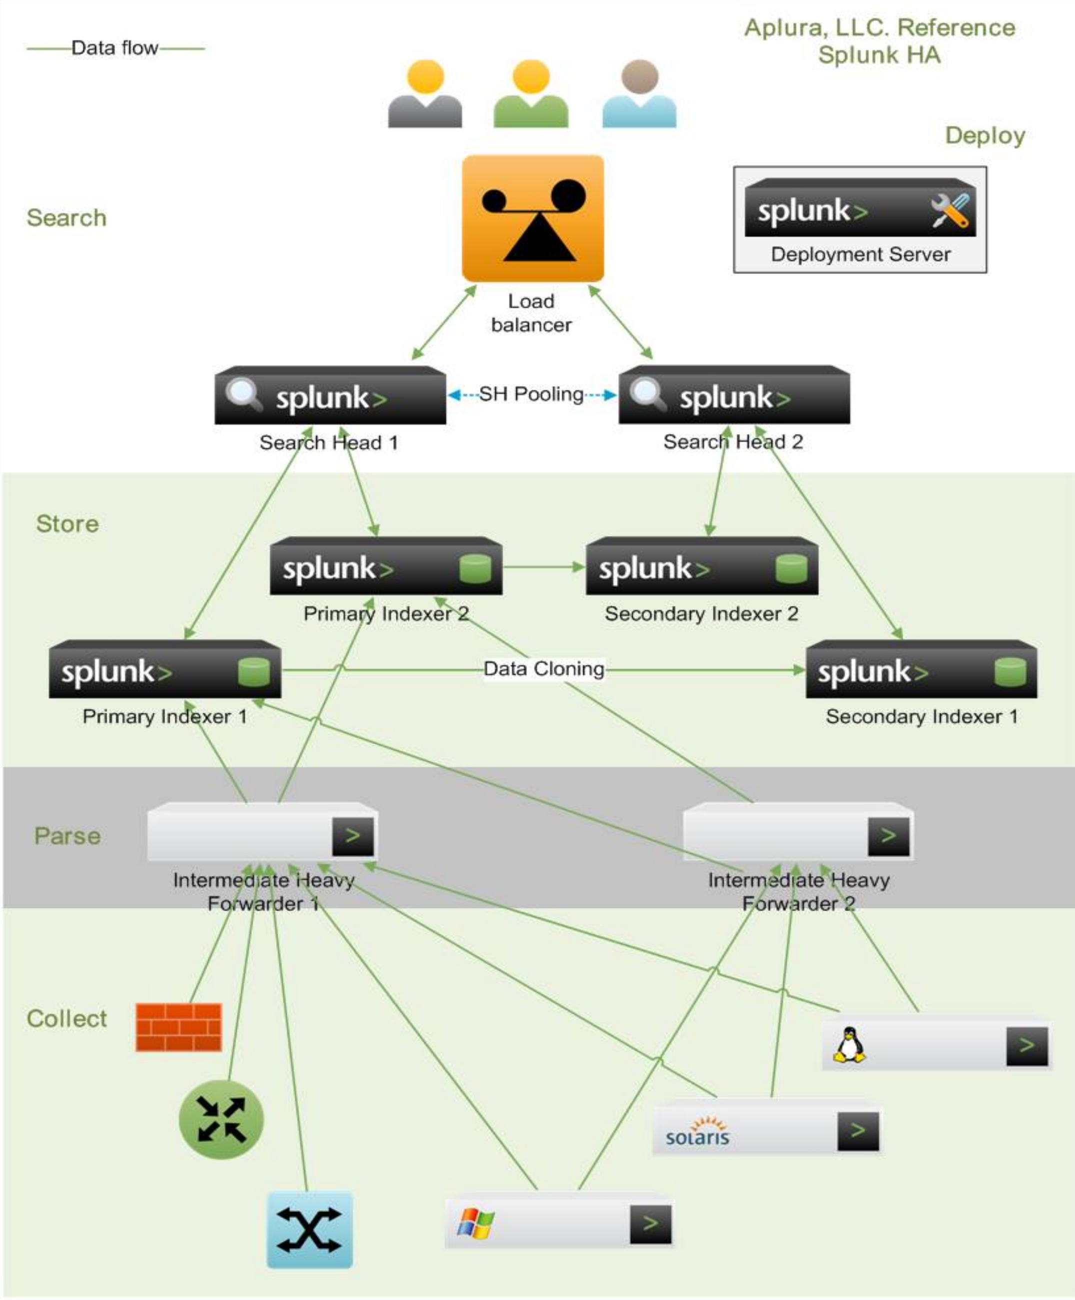

The following are common questions

and problems that frequently arise

during Pro Serv engagements

What does deployment server do?

- Remotely controls Splunk configs for Ent.
  - Uses Splunk apps
  - Runs with whatever rights Splunk has
- Universal (supports whatever Splunk does)
- Can be very powerful and flexible
- Esoteric and docs are generic
- PT: We include it in almost every engagement

Can We Use DS For All Configs?

Yes you can, but is what what you want?

Its a WYSIWYG (or WYSIWYE) world for users.

Wikis, CMS, Google ...

They want to edit what they see in a browser.

PT: Use DS for all non-viewable configs

PT: Beware of local user-changes

PT: Beware of host-specific items

PT: Be Patient (and check your logs)

Do We Need UF, we have ...

Do We Need UF, we have ...

Splunk UF

- Encryption, Compression, Checksum, Load Balances, Auditing, and Store&Forward
- Lightweight to run and manage
  - Doesn't parse data (that is left for the next tier)
- UF free to install and manage for Enterprise

PT: Reduce Service Impact

PT: Change your password

Send Raw Data to Indexers?

You can. Indexers are designed to parse and index the raw data.

It doesn't mean you always should.

PT: In many situations, use an Intermediate Heavy Forwarder

- Brokers TCP sessions from Universal Forwarders
- Parses All Raw
  - Reduces Indexer Load
- Increases Performance
- See BP guide for many other advantages...
- PT: Ensure _internal is forwarded

IHF outputs.conf

```
[tcpout]
forwardedindex.filter.disable = true
forwardedindex.1.blacklist =
indexAndForward = false
maxQueueSize = 1024MB
defaultGroup = index_cluster

[tcpout:index_cluster]
autoLBFrequency=60
autoLB=true
compressed=false
server=sp-indexer1:9997,sp-indexer2:9997
```

We Are Decentralized, Where Do We Place Indexers

- Customers often place Indexers near data
  - They believe this will save on bandwidth
  - This assumption is usually false
- Generally, position indexers near SH
- Exception
  - If users are at both locations, then you have to decide

De-Centralized Example 1: KC/PHX

- All main operations were in Kansas City
- Satellite office in Phoenix with slow link
- Customer placed Indexer in PHX thinking the central SH would just query on occasion
- Result is, all data is queried several times per day.
  - Requires more bandwidth than if SH/INDX were local
  - See next question about data planning

De-Centralized Example 2: NY/UK

- Central IT is in NY, Offices in NY and London
- UK users need only their own data
- NY users need only their own data
- With an Indexer in each location, central IT could search both. This would require maintaining (2) search environments
- Fringe case

We have a **lot** of data. Too much?

- Capacity planning for systems/NICs/bandwidth
- Load-balancer, NIC-teaming, and 10G oh my
- Use numbers, not hunches (when possible)
- PT: Compute per-second incoming average
  - Calculate via Splunk license
  - Calculate via incoming data
- Quick averages don't account for peak/valley

Calculate from Splunk License

- Take a 50GB license
- Divide by daily seconds (86,400)
- 50000000/86400=578 B/s; x 8 to get bits
- 578*8=4624 b/s = 4.6kb/s
- <.5% of a 100baseT or .0005% of a 1G NIC

Calculate from Data

- Ex1: Daily syslog volume = 10G
    - 10000000000 bytes; multiply by 8 and get
    - 80000000000 bits; divide by 86400 and get
    - 925925 bits/s or 925 kb/s
    - <1% of a 100baseT or .001% of a 1G NIC
- Ex2: EPS are 1,200, avg event is 250 bytes
    - 250*8 = 2kbs; multiply by EPS
    - 2400*2kbs = 2400kbs
    - <3% of a 100baseT or .003% of a 1G NIC

How do I search another SH?

- You don't really search the SH directly
- Configure the *outputs.conf* on your SH to send results to the indexers
    - Referred to as “Spraying Data”
- Improve search performance
- PT: Plan for this extra data volume on indexers
- PT: Create all summary indexes on Indexers

How do we test apps?

- Many customers deploy test apps to prod. SH
- New apps may collide with current config
- Unwanted apps may consume resources
- PT: Consider a test/dev Search Head (VM works)
  - Ties in to **production** data without causing harm
  - This is a great place to try out new configs
  - You can also test your own apps here
  - Test upgrades here first

Is it ok to just have one sourcetype?

- Generally, no.
- The field *sourcetype* is special to Splunk
- If you are onboarding all data via syslog (and one sourcetype) consider sourcetype_routing

From *transforms.conf*:

```
[force_sourcetype_for_xxxx]

DEST_KEY = MetaData:Sourcetype

REGEX = write_your_regex

FORMAT = sourcetype::new_sourcetype
```

Is it ok to just have one index?

- Sure, its up to you.
- Index separation allows ACL, Retention, Speed
- If you are onboarding all data into main but want to separate some out, consider index_routing

From *transforms.conf*:

```
[force_index_route_for_xxxx]

DEST_KEY = _MetaData:Index

REGEX = write_your_regex

FORMAT = new_index
```

Index/Sourcetype Routing

- Index/Sourcetype routing happens at parsing
- If IHF are in use, then apply the config there
- If no IHF, then usually apply on the indexers

When do we create a new index?

My 2.5 Questions Rule for Creating a new Index:

When do we create a new index?

My 2.5 Questions Rule for Creating a new Index:

1. Data has a different retention policy? (Y/N)

When do we create a new index?

My 2.5 Questions Rule for Creating a new Index:

1. Data has a different retention policy? (Y/N)
2. Data has different access restrictions? (Y/N)

When do we create a new index?

My 2.5 Questions Rule for Creating a new Index:

1.Data has a different retention policy? (Y/N)

2.Data has different access restrictions? (Y/N)

2.5 You want to specifically include/exclude data e.g. index=foo or NOT index=foo? (Y/N)

When do we create a new index?

My 2.5 Questions Rule for Creating a new Index:

1.Data has a different retention policy? (Y/N)

2.Data has different access restrictions? (Y/N)

2.5 You want to specifically include/exclude data e.g. index=foo or NOT index=foo? (Y/N)

If you answered “Y” at all, then create a new index

How do we know the data is right?

Validate

How do we know the data is right?

Validate

VALIDATE

How do we know the data is right?

Validate

VALIDATE

V – A – L – I – D – A – T – E

How do we know the data is right?

Validate

VALIDATE

V – A – L – I – D – A – T – E

VALIDATE

Validate – 10 Min. Time Audit

- Time is broken “everywhere”
- But we have NTP...
- IM/IR needs reliable time
- PT: Perform a monthly time-audit (schedule it)

Validate – 10 Min. Time Audit

- Time is broken “everywhere”
- But we have NTP...
- IM/IR needs reliable time
- PT: Perform a monthly time-audit (schedule it)
  - Splunk Search → sourcetype-by-sourcetype
  - From timepicker select “Realtime/All-Time”
  - Watch “Timeline” for visual queues

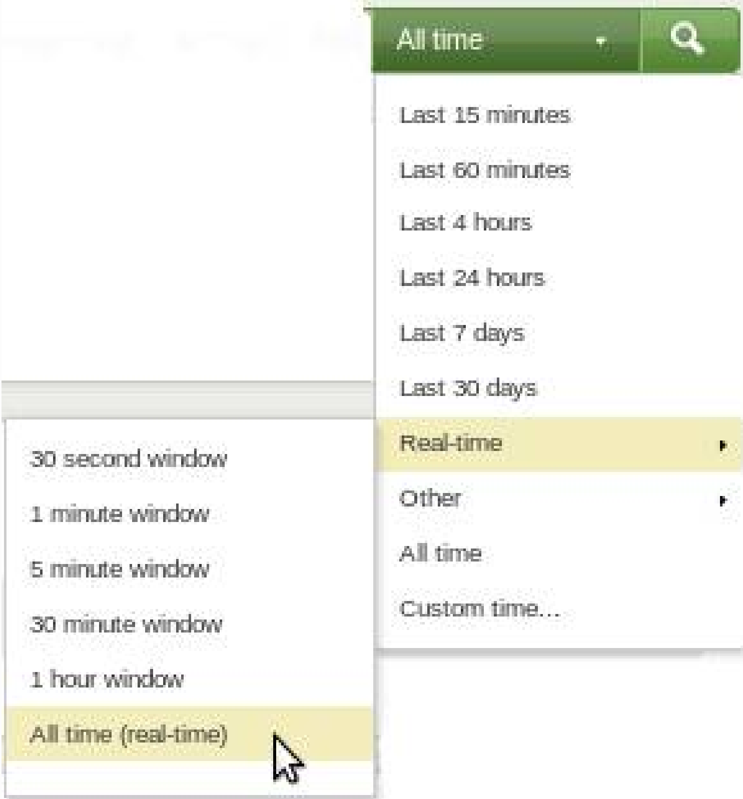

Validate – Duplicate Data

- Double+ data exists at almost every customer
- Consumes license, consumes storage
- Inaccurate **count** reporting
- Don't monitor both the raw syslog and archives
- Don't use syslog and a UF on the same host
- PT: Query for double-data on occasion

Validate – Missing Data

- Too often customers are missing valuable data
- You discover when you need it...too late
- Deployment Monitor helps but isn't perfect
- PT: Add alert after onboarding a new data type

Can storage be added later?

- Yes, you can change almost anything you want
- Leverage tiers effectively (hot, warm, cold)
- Adding data might mean data migration
  - Data migration can be tricky
- PT: Get Storage Folks Involved Early!
- PT: Configure data retention before it becomes an emergency

TA-uas_parser

- User-agents are difficult
- There are thousands of them
- Parsing them is a nightmare

- People are already doing this: user-agent-string.info
- Uses their library to enrich User-agent information with easier to use fields

Domain Categories

- Uses data from urlblacklist.com
- Not just blacklisting, but also categories for domains
- Category information for more than 3.5 million domains
- Data at $6 per month

- Can be applied to any events with domain name information
- Web Proxies
- Email
- DNS

Content Available Now!

- Talk: www.aplura.com/splunklivedc
- BP: www.aplura.com/splunkbp

APLURA

FOCUSED INFORMATION SECURITY