Antipatterns
It seemed like a good idea at the time…

Duane Waddle

Senior Security Engineer, Defense Point Security

David Paper

Senior Advisory Engineer, Splunk

.conf2016: YOU make it great!

Disclaimer

During the course of this presentation, we may make forward looking statements regarding future events or the expected performance of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward looking statements we may make. In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only and shall not, be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionality described or to include any such feature or functionality in a future release.
Godfrey for President in 2016!

Patterns Vs Antipatterns

- Design Pattern: A good, repeatable approach
- Anti-Pattern: Looks like it could be a pattern, but is the opposite
- Potentially Inflammatory - YOU MAY NOT AGREE WITH US
- Development world: design patterns evolve, mature, decay

- Questions? Throw ‘em out during the talk
- Discuss the merits? Hold ‘til the end & potentially over refreshments

Who Are We?

- Duane
  - Recovering sysadmin
  - Enjoys Bourbon
  - Proserv Consultant, has seen nightmares
  - Slack & IRC #splunk: duckfez
- Dave
  - Rose colored glasses wearing former Ops guy
  - Didn't set out to experience ***every*** antipattern
  - Learn, document, share, repeat
  - Slack & IRC #splunk: cerby

Intermediate Forwarders

- UF UF UF
- Keep funnels at bay
- Maintain 2:1 pipeline ratio
- Adjust autoLB freq down, 5-10s
- DDoS yourself
- Bandwidth / StS protocol considerations

Simplified IHF Funneling Example

- 1000 UFs, 10 indexers, 10 EPS per UF
  - Each indexer sees (statistically) 100 UFs, 1000 total EPS
  - Statistical variances at the indexer are minor and unremarkable
    - 95 connections vs 105 does not make a huge EPS difference

Simplified IHF Funneling Example

- Now let's add 5 Intermediate HFs
  - 1000 UF => 5 HF => 10 Indexers
  - Each HF now has (statistically) 200 connections, 2000 EPS
  - 2000 EPS from each HF is now passed on to 1 indexer
  - Statistical variance in the HF => IDX connections make BIG difference
    - Some get 0 EPS, some get 4000, maybe 10,000 in worst case!
  - Cooked?
  - How many pipelines

Virtualization

- Multiple indexers on single ESX = Bad
- IO funnel to storage
- Smooth out peaks/valleys? M/R load hits at the same time
- CPU Cores: Physical/virtual, reserved
- Hops to Disk: 3 BM, 7 VM to SAN
- VM over subscription: Storage Admins & Perf Tuners
- Resource validation on day 1, 30, 90, 180, 365?

Hardware Specs

- More nodes = Faster individual search
- More cores = Higher concurrency
- Friendships
  - # of cores = Concurrency
  - Ghz = Throughput
  - RAM ≈ IOPS
- LB gets in the way, unsupported
  - Single dest breaks FWD balancing logic

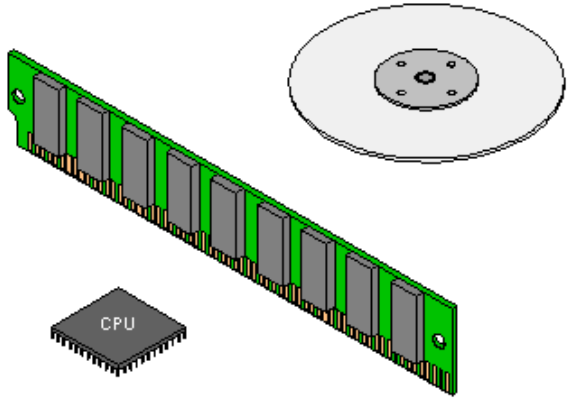

Bacon Flavored IOPS

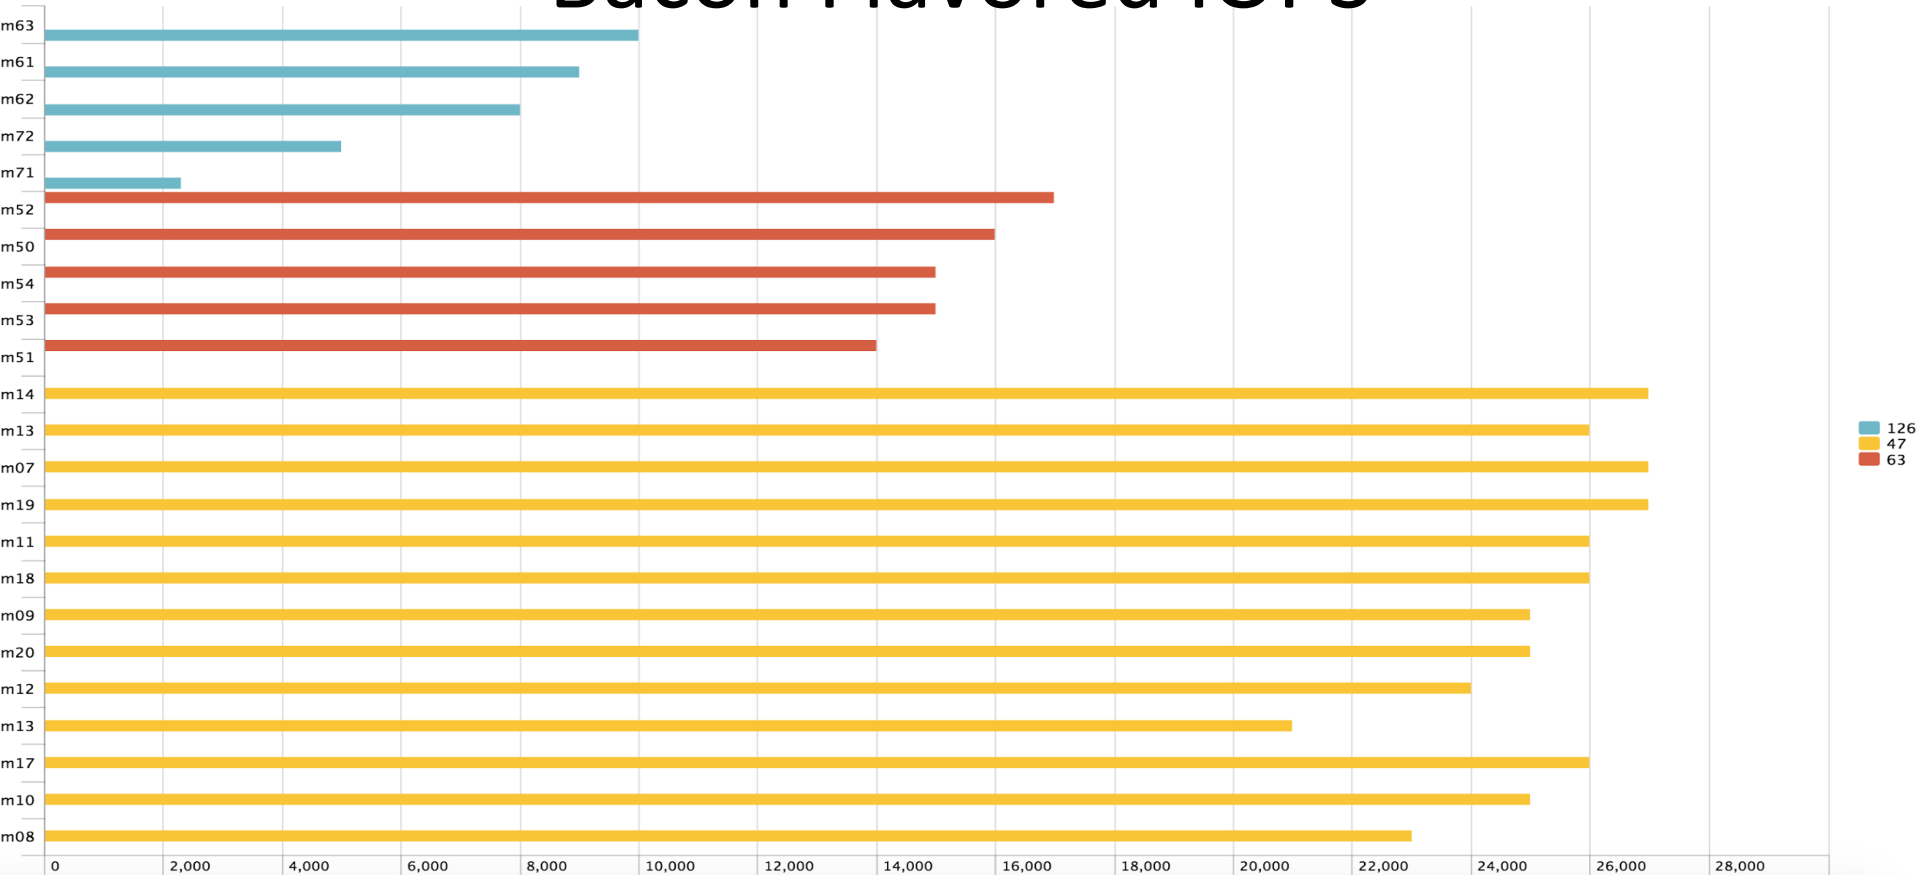

Real Time Alerts

- RT sells monitoring software
- Who's paying attention 24/7?
- Automated remediation tied to RT alerts? Mature environment!
- Email or NOC notification? RT unnecessary
  - 1, 5 or 15 minute scheduled alerts
- Alert fatigue

Index Sprawl

- Env “of size”: < 100: too few; > 1000 too many
- Triumvirate of Access Control, Retention, Performance
- Inode exhaustion, file handle explosions
- More work for warm -> cold migrations
- Clustering limiting factor: Too many buckets

We Don’t Need No Stinkin’ Root

- Can/should run as non-priv user
- Remote code exec
  - NOT A SUBSTITUTE FOR REAL CONFIG MANAGEMENT / SOFTWARE DIST TOOLS
- Friends in low places: FS Perms & ACLs
- Ports > 1024
- Bind to 0.0.0.0. Ponder implications
- See Matt Uebels talk for more details

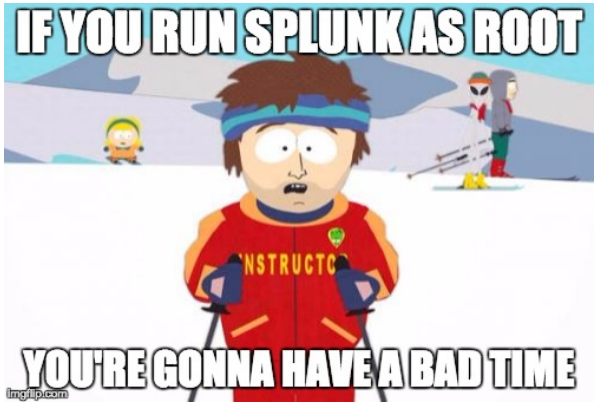

Fidelity Of Data Matters

- Low fidelity Data Collection (Syslog) robs the future
  - UF > SNARE, 101% of the time.
- Specific formats expected
  - Altering nukes prepackaged extractions (ES)
  - access_combined is a good example - add K=V to the END of the event to not have borked extractions

FIN

Bring It

THANK YOU

splunk>