A NEW PLATFORM FOR A NEW ERA

Unveiling Clusters of Events for Alert and Incident Management in Large-Scale Enterprise IT

Derek Lin*, Rashmi Raghu, Vivek Ramamurthy, Jin Yu, Regunathan Radhakrishnan – Pivotal
Joseph Fernandez – Visa Inc.

KDD 2014, 8/26/2014

Pivotal

Primary Motions for Pivotal

- *Agile:* data driven apps and rapid time to value
- *Data Lake:* store everything, analyze anything
- *Enterprise PaaS:* revolutionary software development and speed; build the right thing

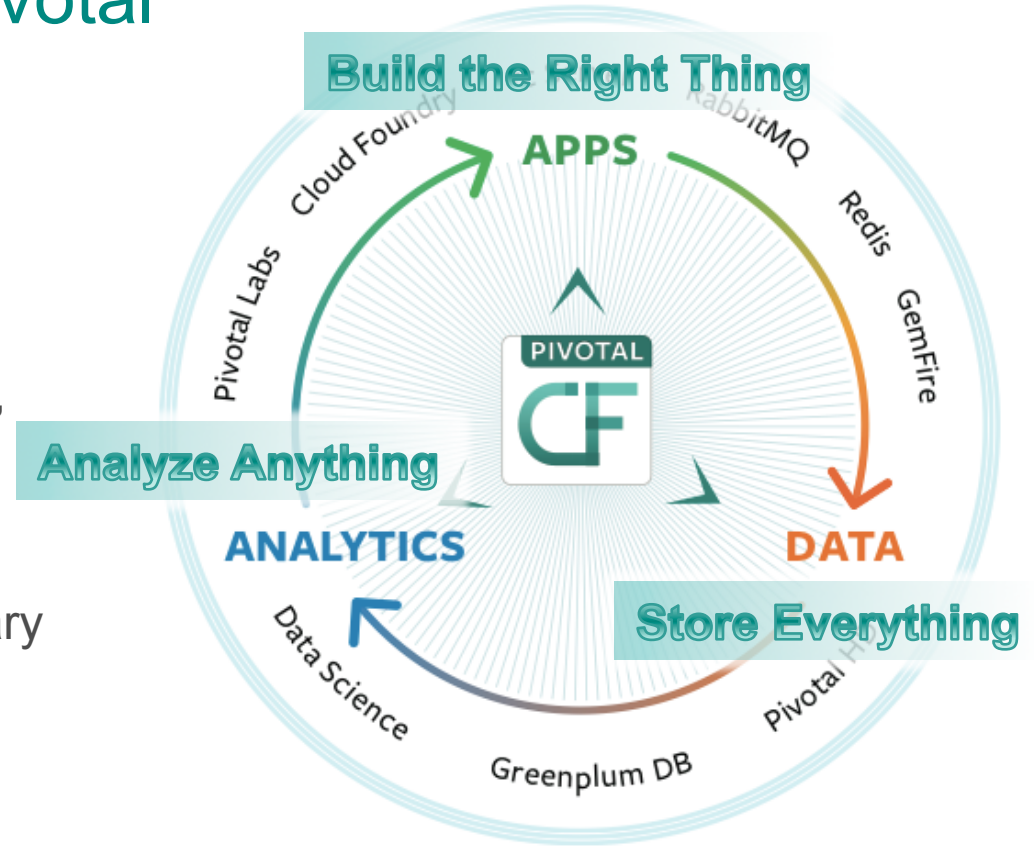

Pivotal

Motivation

- Enterprise network is complex
  - Different technology components with lots of dependency
- Role of Reliability Engineering
  - Ensure 24x7 uptime, network monitoring and quick resolution
- Alerts are high volume; eyes-on-glass operation
  - Event logs, performance metric log, incident tickets
- What intelligence can we mine from data to improve operational efficiency for Reliability Engineering?

Pivotal

IT Infrastructure and Data Sources

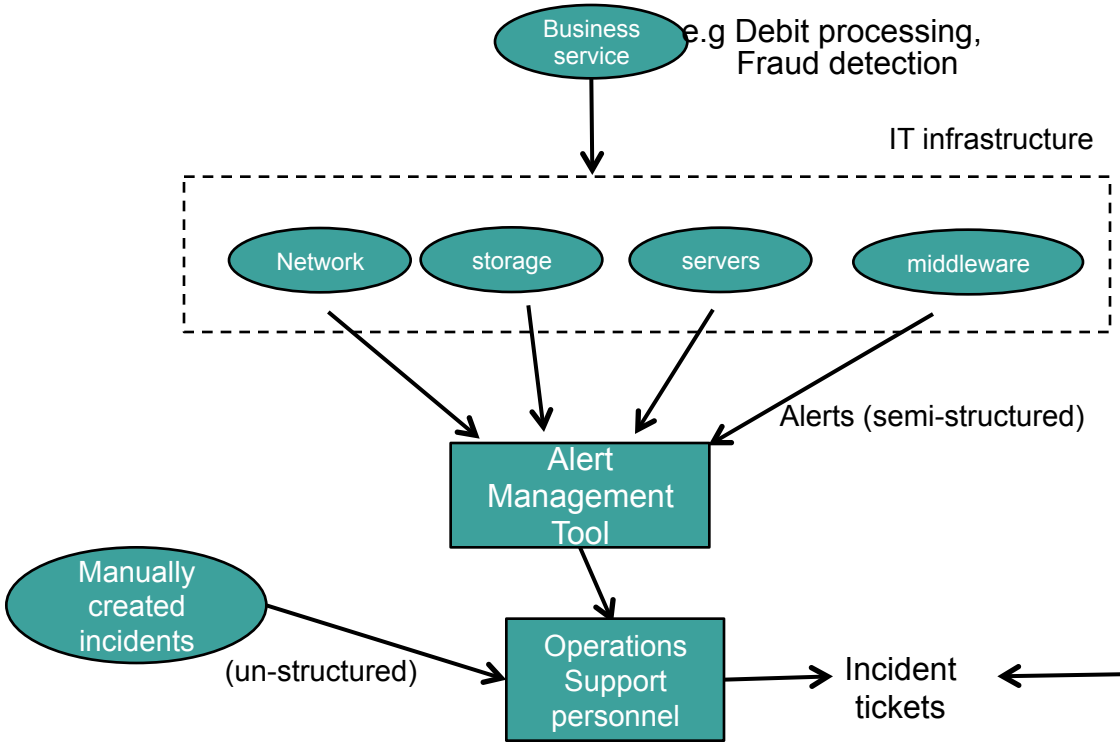

- Alerts
- Incidents
- Incident work info

Business Goal

- Improve operational efficiency of IT Support & Help Desk
- Method: Analyze and cluster historical alert and incident data :
  - To profile clusters in volume or mean-time-to-resolve
  - To recommend incident resolution
  - To predict infrastructure failures before they occur

Challenges

- Data
  - Large volume: 10 million alerts and incidents in 6 months from just one business service. There are numerous services – debit processing, fraud detection, etc..
  - Multi-structured: Semi-structured and unstructured text
  - No labeled data
- Analytics
  - Alerts/incidents have short text
  - Clustering techniques at scale for incidents
  - Cluster interpretability for qualitative evaluation

Pivotal

Approaches

- Unsupervised learning based approach
- Represent alerts as bag-of-words; Define a distance metric to compare any two alerts and perform clustering.

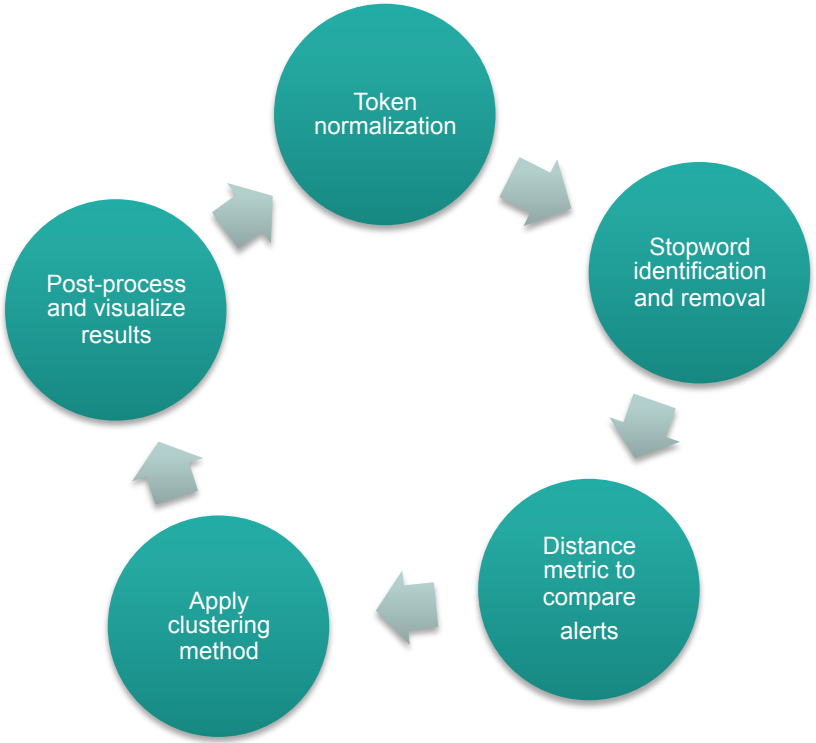

Pivotal

Clustering of Semi-Structured Text Alerts

- Distance metric
  - Jaccard Index, dist(A,B) = 1 – (|A ∩ B|/ |A U B|)
- Given NxN, create a graph, establish an edge between nodes (alerts) if the distance < h
- Clustering
  - Connected components detection
  - Graph partitioning

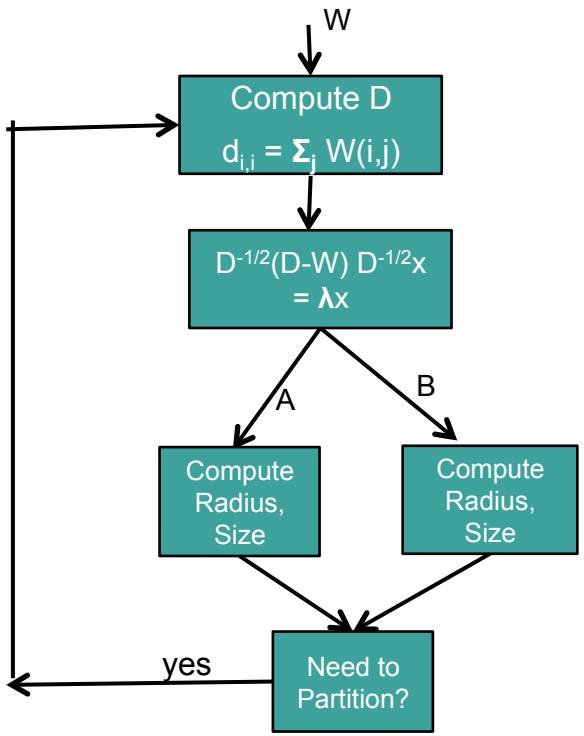

Clusters From Each Alert Group

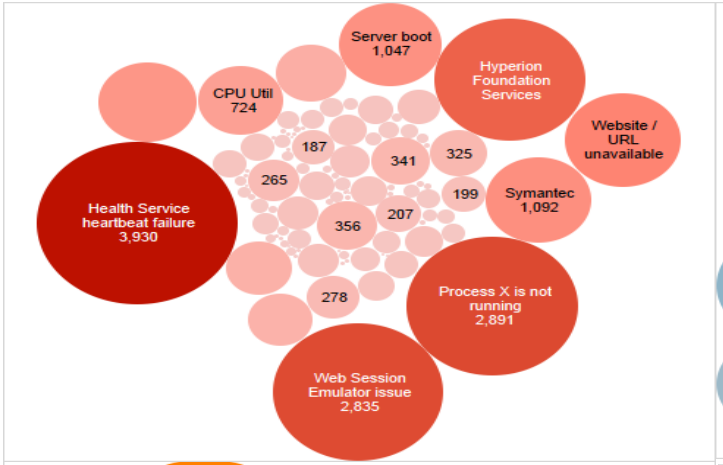

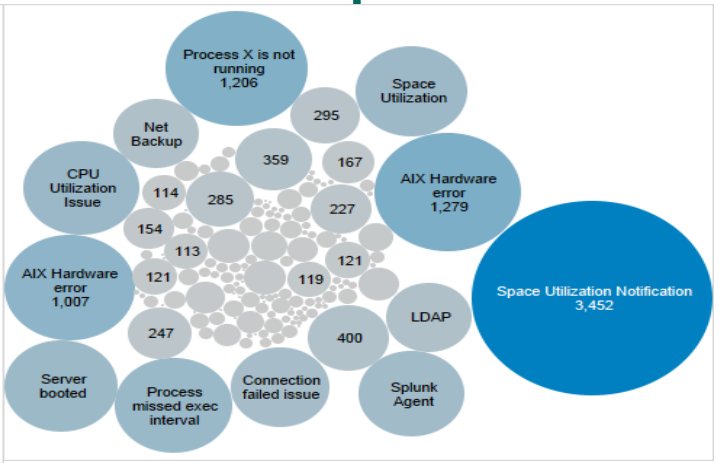

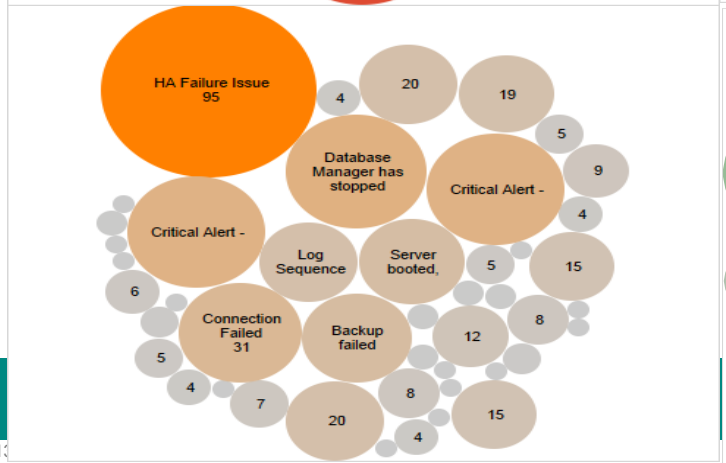

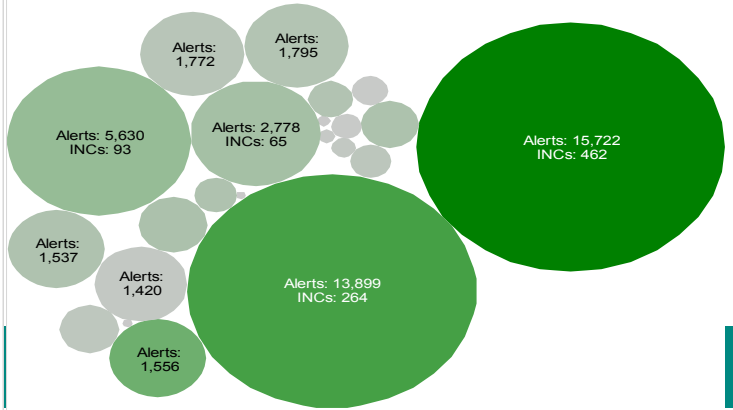

Pivotal

Mean-Time-To-Resolve for Clusters

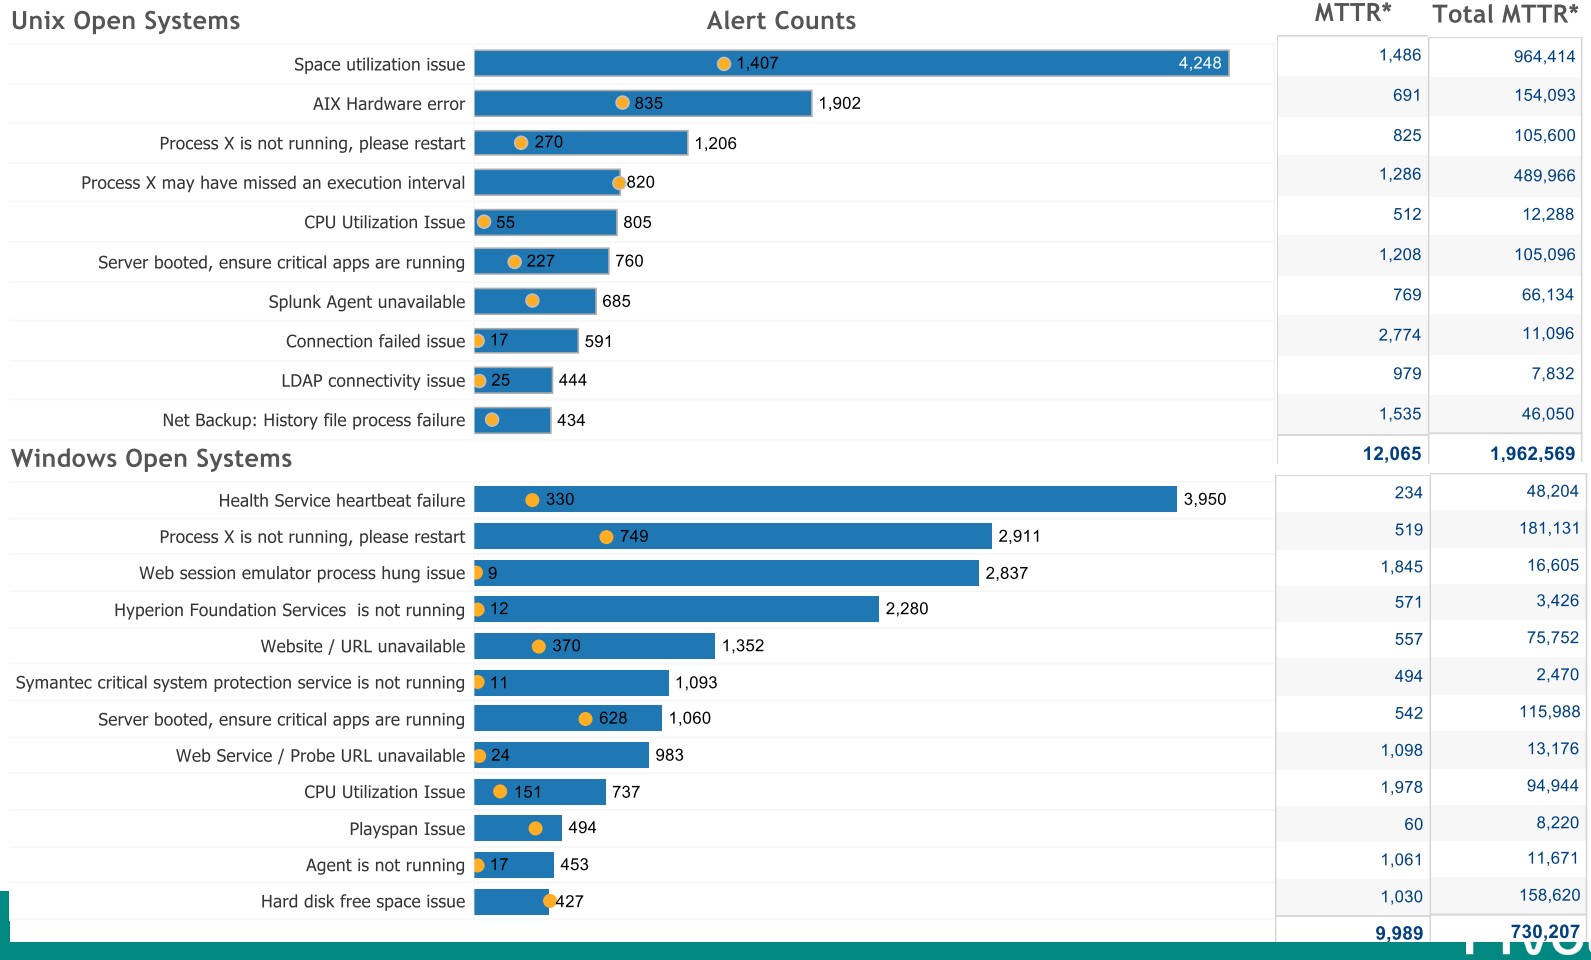

Cluster Visualization – An Alternative To Word Clouds

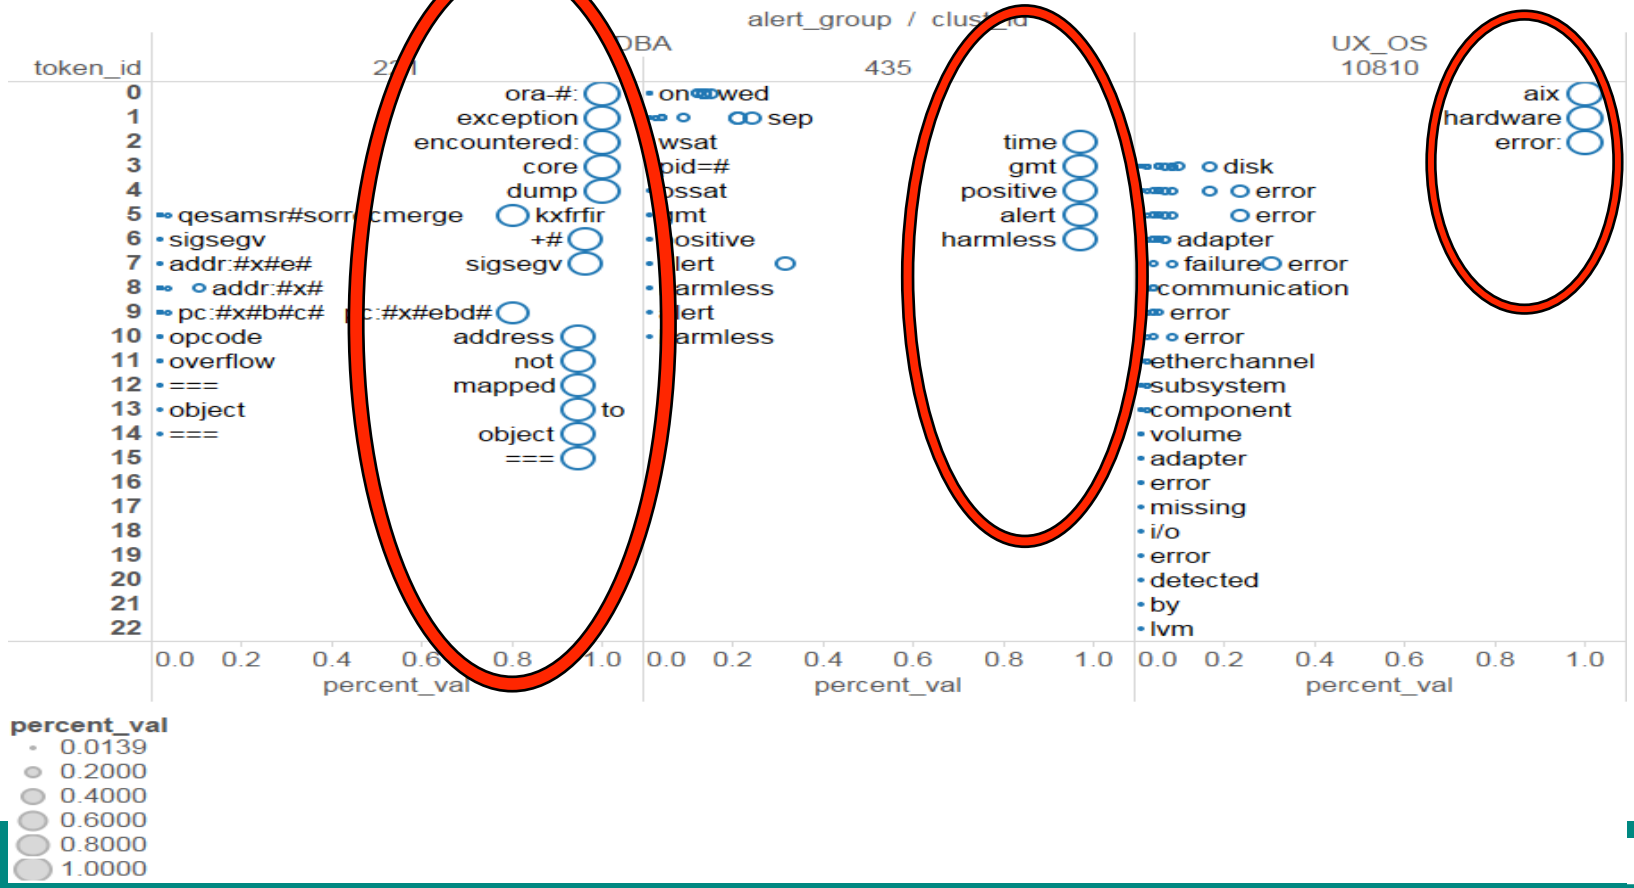

Clustering Unstructured Text Incidents

- Distance metric
  - Dist(A,B)= Const - |A ∩ B|
- Clustering
  - Use of non-negative matrix factorization to find latent representation of incidents
  - KD tree for coarse space partitioning
  - Hierarchical clustering: complete-linkage
  - Refine by merging clusters
  - Choose a prototype for cluster representation
    - An examplar with the shortest distance to its farthest member in a cluster

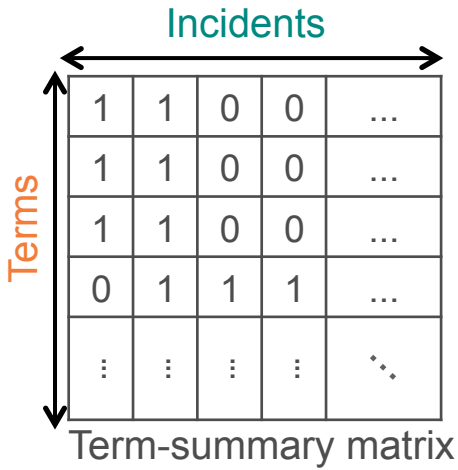

Top-25 clusters account for 20% of 67K incidents

cluster_id	count	cumulative_pct	keywords	exemplar_summary
1	2400	3.59%	ended,notok	VSA1043E DE04 3 Q622A551 134 ENDED NOTOK 20130721 23:06:29 1064 MVS
2	1892	6.42%	alert,splunk	mibaker-4-sl55pvimrep03-splunk alert: The fileextract-primary process is not running.
3	925	7.81%	percent,utilization	cayuban-3-sa73cmlsp02-/var is at 89.93 percent utilization
4	758	8.94%	booted,ensure,running	mibaker-4-sw730ferispa01-SCOM:Server booted at 2013-09-27T21:59:41.0000000 00:00 - Please ensure critical applications are running.
5	567	9.79%	provisioning,vblock	Automated Vblock server provisioning request for sl73vcasapq001 [VCAS Files Processing System]
6	527	10.58%	address,email	Outlook: DIfferent email address.
7	505	11.34%	password,reset	Please reset the password for ROETLQ on SA73ROLDBQ51 server
8	464	12.03%	aix,hardware	dcanong-3-sa73vctr02-AIX HARDWARE ERROR: DISK OPERATION ERROR, CONNECTION FAILURE, UNABLE TO COMMUNICATE WITH DEVICE, PHYSICAL VOLUME DEFINED AS MISSING, I/O ERROR DETECTED BY LVM
9	442	12.69%	distribution,list	SSSAdmin: Distribution List Request/Change/Deletion - NA - QA (DCM)
10	400	13.29%	supplier	Supplier Setup - My issue is not listedremove end date for employeeSupplier name: BOI, SHOOK WAISupplier number: 14009Inactive on: 05-APR-2013Supplier site name: OFFICEInactive on: 05-APR-2013

Pivotal

Takeaways And Lesson Learned

- Understand where your alerts/incidents are coming from is an important step in improving infrastructure support
  - Profiling classes of alerts for business intelligence
  - Resolution recommendation
  - Application or hardware failure prediction
- MPP computing architecture + proper algorithm choice are needed to deal with scalability
- Tuning iterations require good cluster visualization
- Future direction: leverage temporal info.

Pivotal

A NEW PLATFORM FOR A NEW ERA