An Enterprise Dynamic Thresholding System

Mazda Marvasti, Arnak Poghosyan, Ashot Harutyunyan, and Naira Grigoryan

Presented by: Bob Patten, VMware

Agenda

- Modern IT management challenges
- The data agnostic method for anomaly detection
- An enterprise dynamic thresholding system
- Data categorization approach
- Category-specific dynamic threshold determination
- Experimental insights
- Real-World Customer Use Case

Modern IT management challenges

- Management based on operator domain expertise is no longer efficient
- Huge distributed cloud systems, virtualized environments
- Complicated interrelation between the constituent components
- Business infrastructures are highly dynamic – behaviors do not fit classical Gaussian normal distributions
- Static thresholding of processes and performance indicators become inadequate yielding thousands of un-actionable alerts
- Manually developed and maintained correlation rules yield no significant benefit to problem identification

The data agnostic method for anomaly detection

- Automatically learns the normal behavior of **any** time-series metric
- Makes no assumptions as to the metrics' behavior or distribution
- Calculates an upper and lower bound hourly dynamic threshold
- Determines normal or abnormal behavior (anomalies) of individual metrics
- When metrics behave abnormally, additional algorithms and deterministic methods can be applied to determine system abnormality
- VMware's vCenter Operations Manager (vC Ops) is an industry-leading Big Data solution for IT Management which utilizes the described enterprise dynamic thresholding algorithms

An enterprise dynamic thresholding system

- The monitoring and alerting based on data analysis behind vC Ops

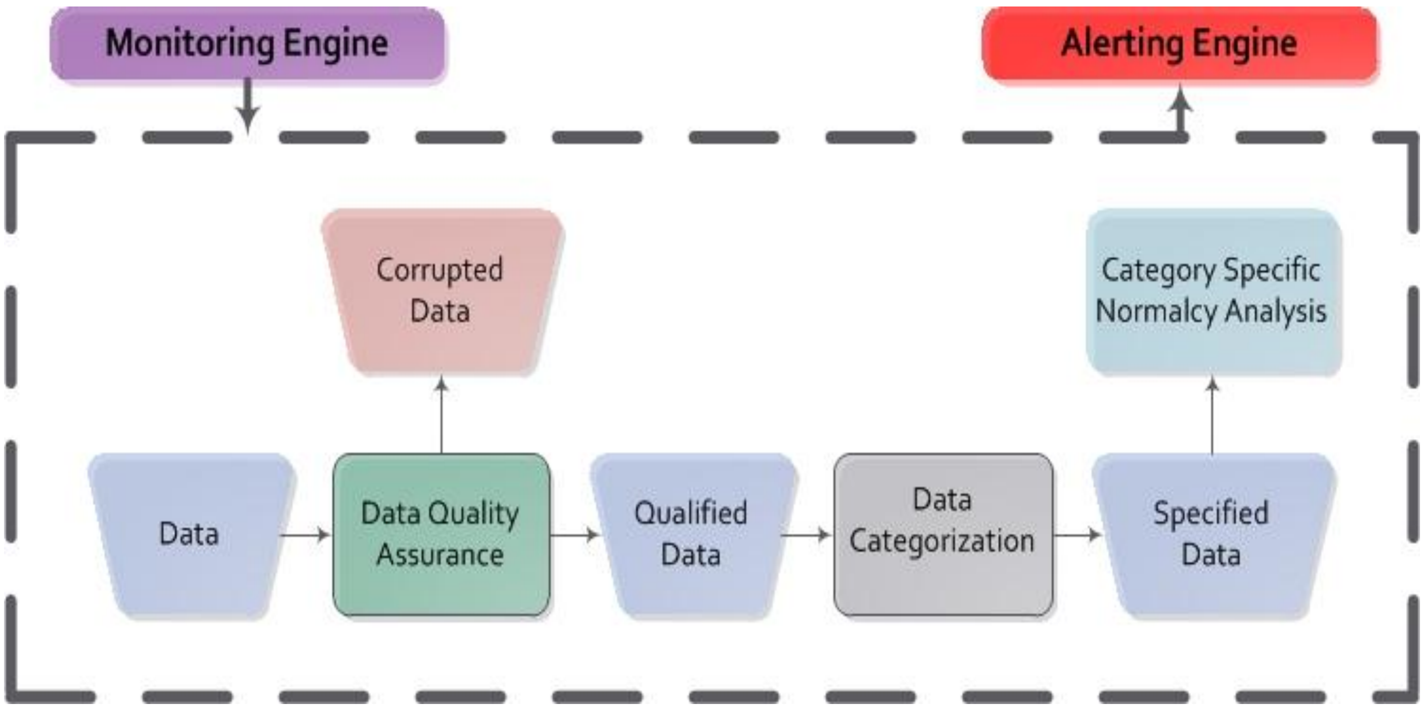

Example metric dynamic threshold

Weekend/Weekday repeating pattern of normal behavior

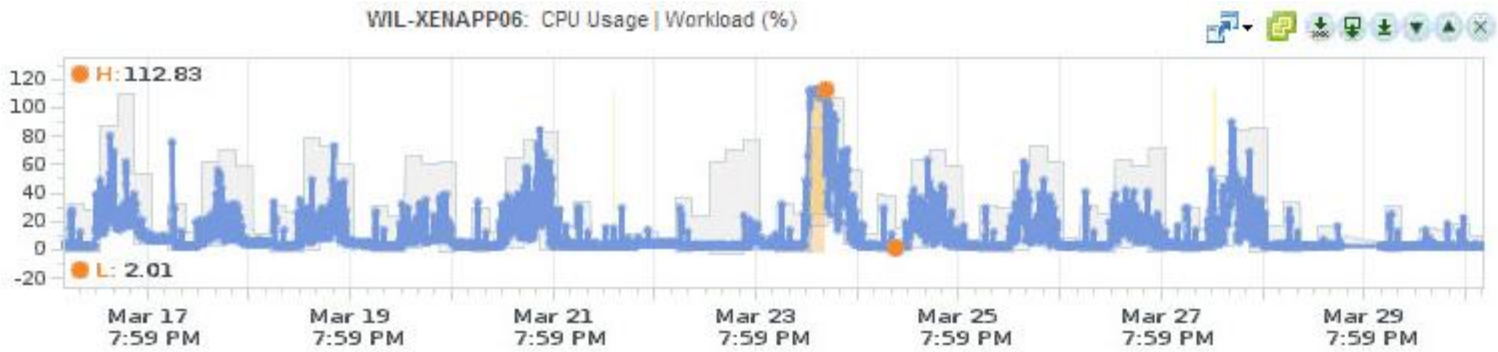

Resulting Dynamic Thresholds

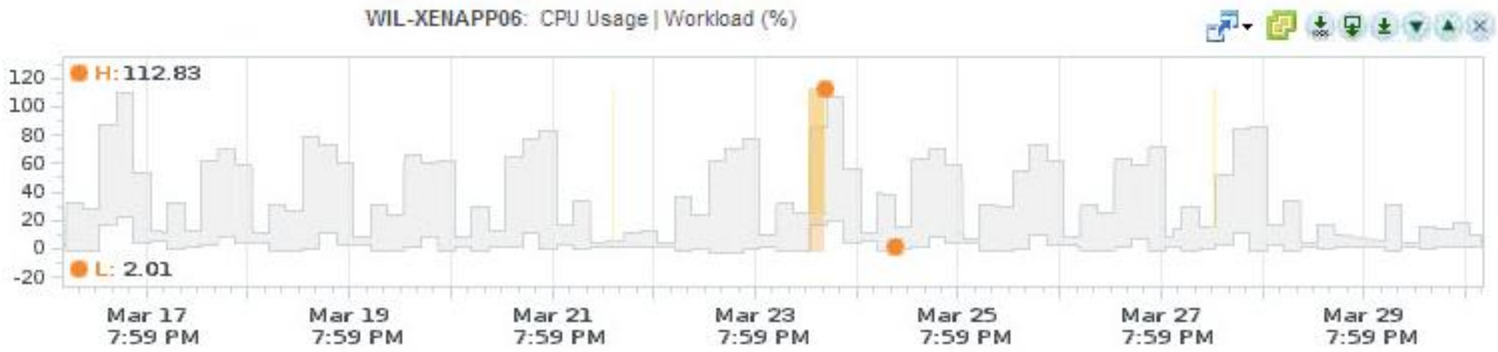

Data categorization approach

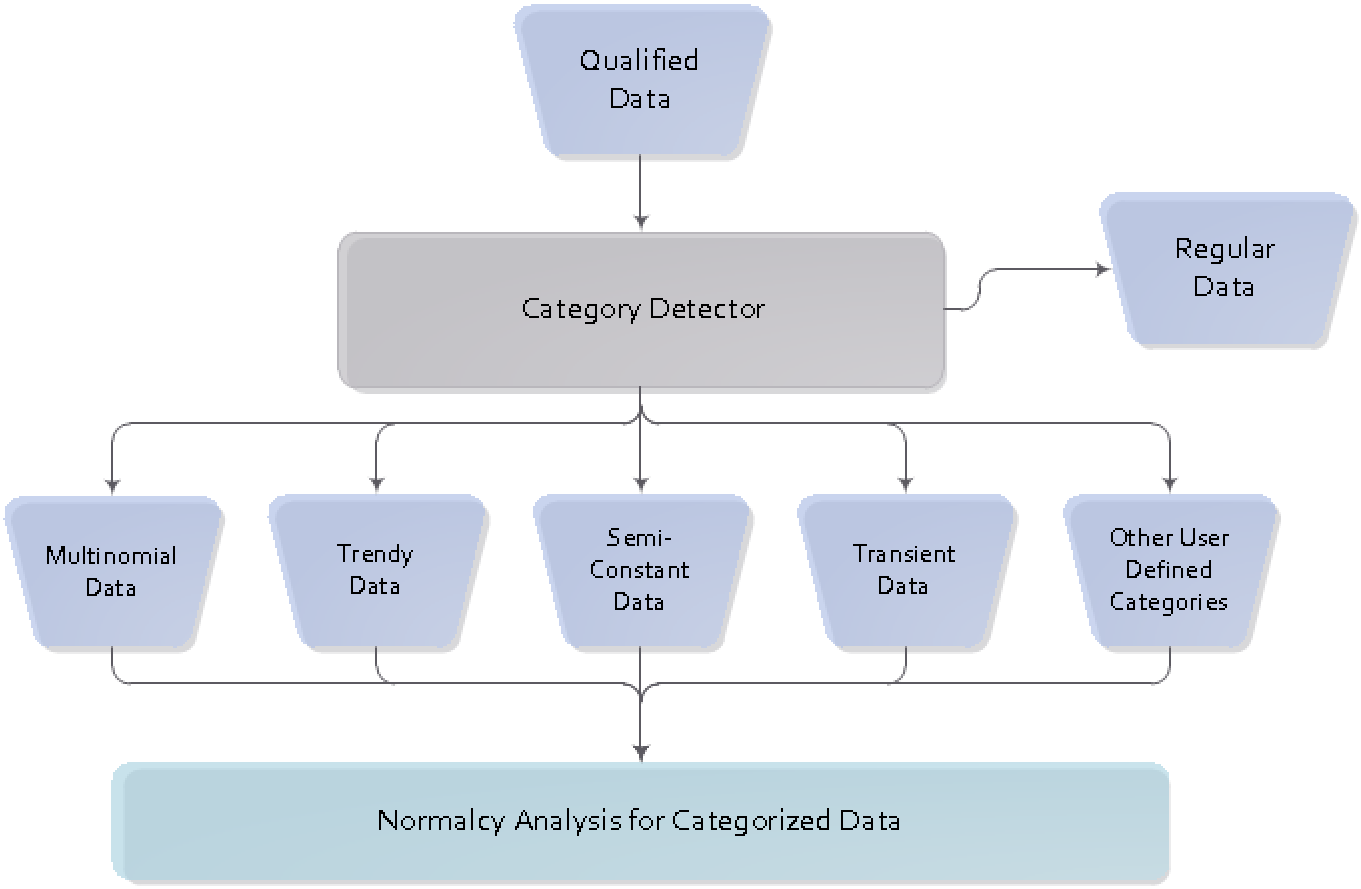

Data categorization approach: examples

- **Trendy**

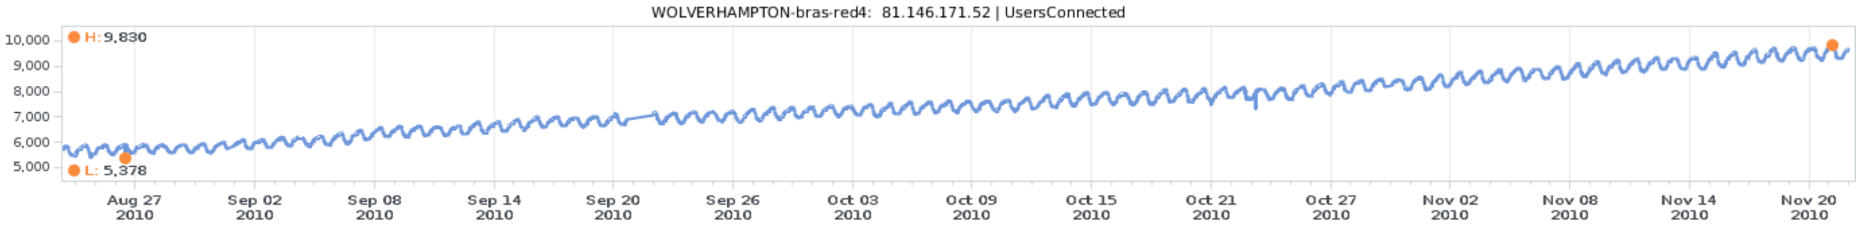

- **Multinomial**

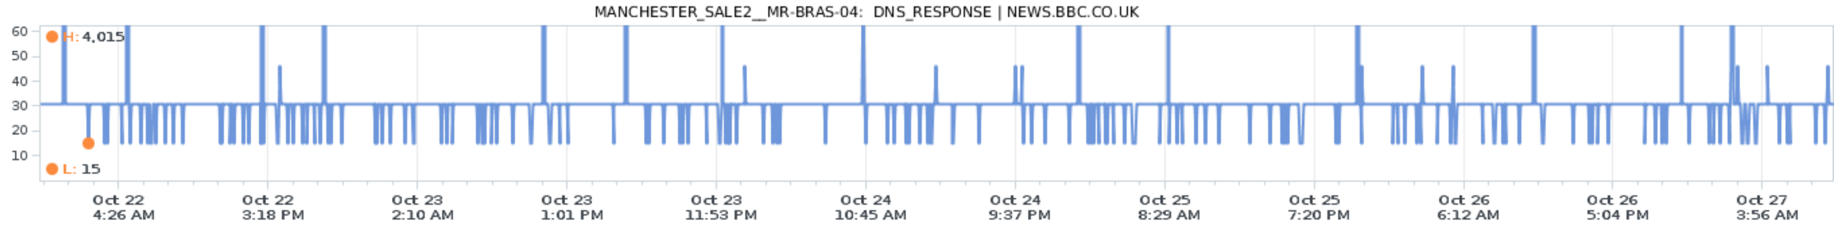

Data categorization approach

- **Sparse**

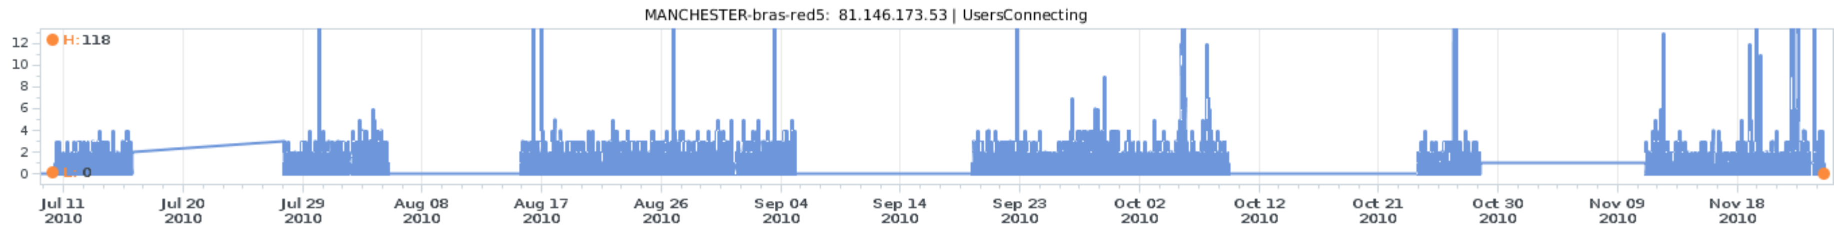

- **Regular/Periodic**

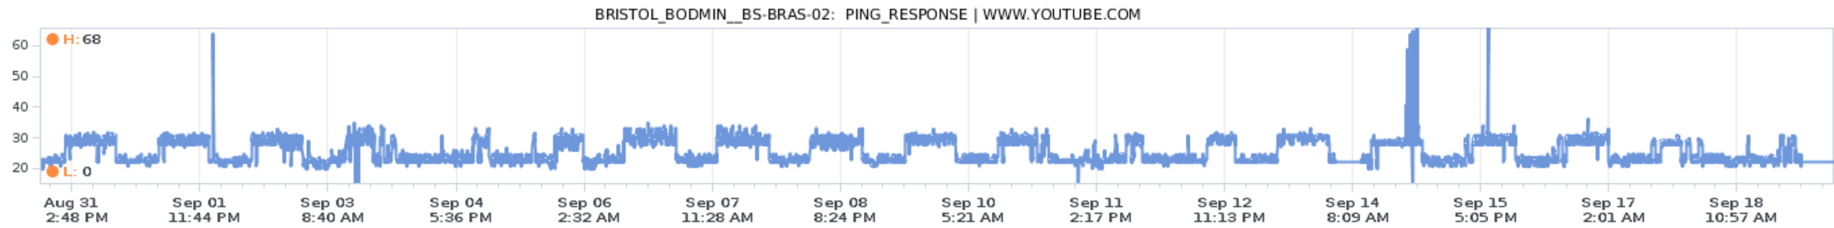

Category-specific DT determination: sparse data

- Performing data density recognition based on probability calculation that reveals distribution of gaps
  - Random?
  - Uniform?
  - Pattern?

- Differentiating the following clusters of data:
  - Data Identification: Dense/Sparse (relative to monitoring interval)
  - Data with technical gap (localized gap due to malfunction of monitoring device)
  - Corrupted Data

Category-specific DT determination: stable data

- Statistical stability recognition of data
  - If data is stable or its stable portion can be selected then the data is defined as **Stable Data**
  - Otherwise data is defined as **Corrupted**

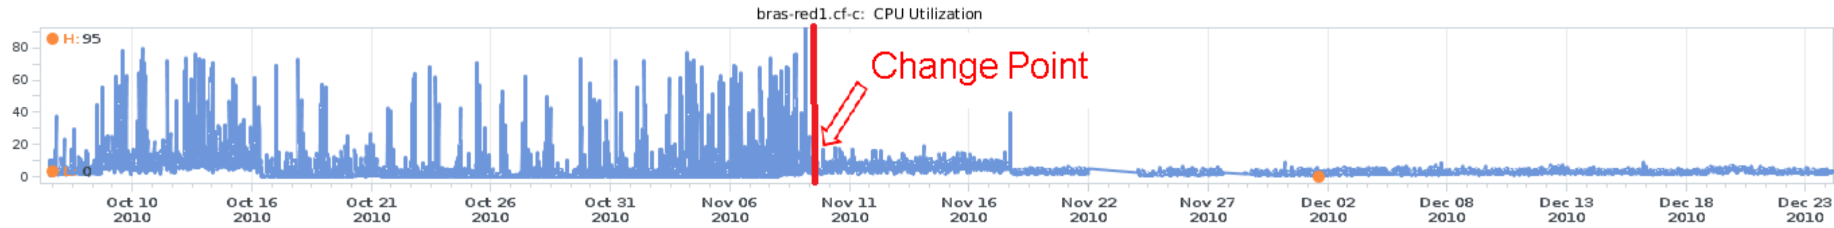

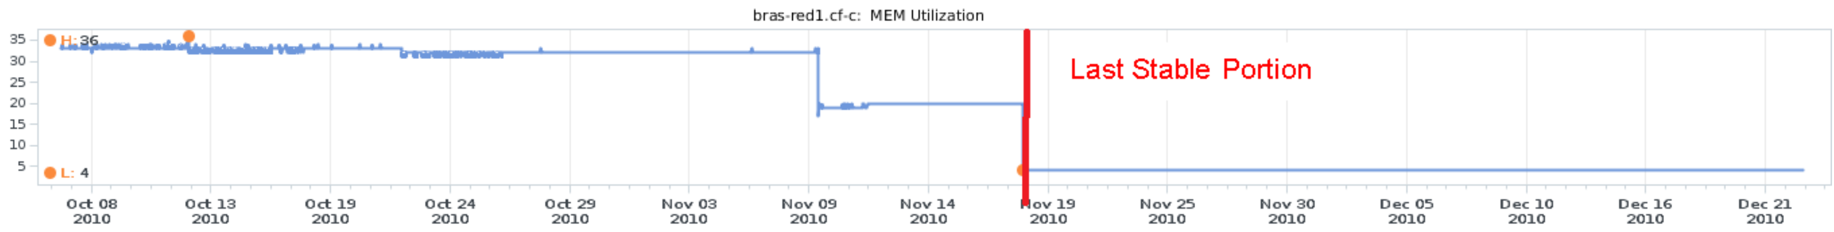

Category-specific DT determination: variability

- $R = \dfrac{iqr\left(\{x'_k\}_{k=1}^{N-1}\right)}{iqr\left(\{x_k\}_{k=1}^{N}\right)} 100\%, \quad iqr\left(\{x_k\}_{k=1}^{N}\right) \neq 0$

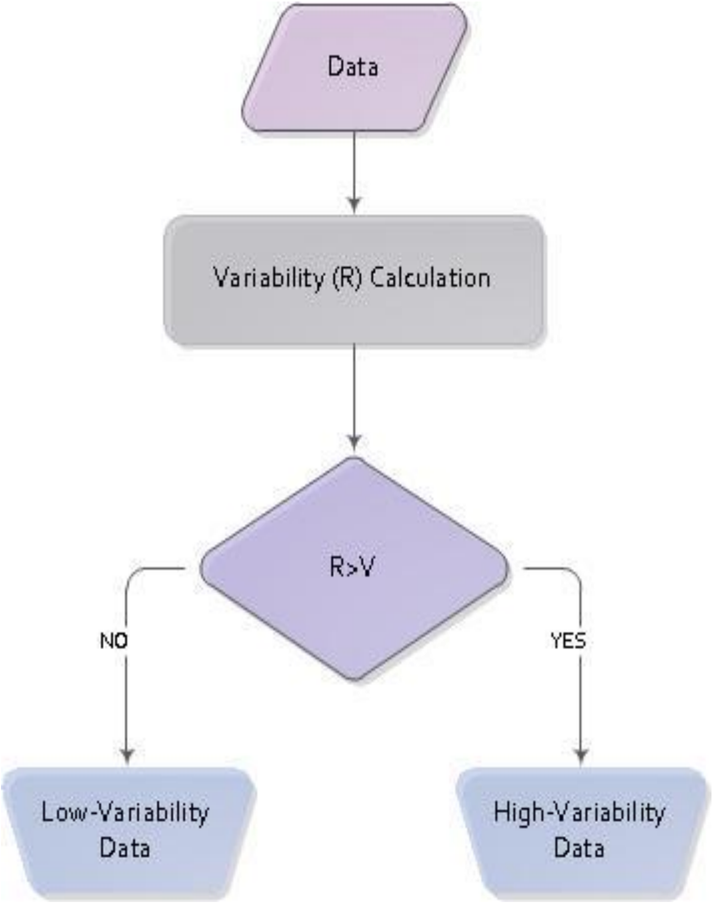

Category-specific DT determination: periodicity

- **Periodic data**: seeking similar patterns in the historical behavior of time series
  - The notion of the Cyclochart is similar to the frequency spectrum in the Fourier analysis or signal processing

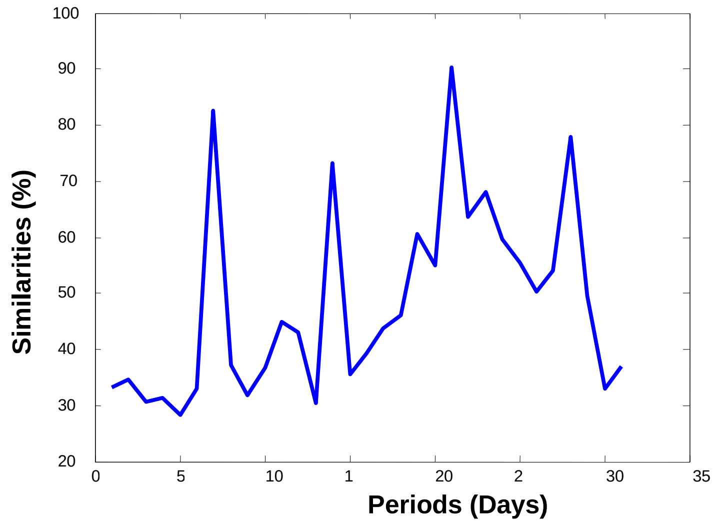

Category-specific DT determination: optimization

- Statistically trade-off the number of false positive and false negative alerts
- Two different approaches for determination of DT's via maximization of the objective function

$$g(P,S) = e^{aP}\frac{S}{S_{max}}$$

  - Data-range-based analysis
  - Data-variability-based analysis

Experimental insights

- A specific customer metric data set
- Selected 3215 monitored metrics
- Those metrics represented the essential flows for one of the customer’s critical business services
- Data length is one month
- Ran metrics through Dynamic Thresholding analytics process
- Resulting count of periodic/non-periodic/corrupted data

Periodic	Non-Periodic	Corrupted	Overall
1511	1595	109	3215

Experimental insights

- Distribution across the categories

Data Category	Count (Percentage) of Metrics in the Category
Multinomial	724 (22.5%)
Trendy	165 (5.1%)
Semi-Constant	532 (16.5%)
Transient	102 (3.2%)
Sparse	88 (2.7%)
Low-Variability	826 (25.7%)
High-Variability	669 (20.8%)
Corrupted	109 (3.4%)

A Production Use Case – 4 Hour Proactive Notification

Production Scenario

- Citrix Xen Desktop Remote Desktop Environment on Virtual Infrastructure
- Multiple XenApp Server VMs serve the end-users Remote Desktops
- Monday morning, March 24th, significant abnormal behavior
- XenApp VM owner (Citrix Admin) called at 8:00 AM, returned call at 10:00 AM
  - Initial evaluation by Citrix admin is **“All OK, end users are not complaining”**
  - Subsequent investigation yielded a call-back and thank you to Operations
    - A config change in the Citrix env over the weekend was causing orphaned sessions
    - Citrix Admin fixed the error and cleaned up the sessions
    - **If Operations had not proactively notified Citrix Admin, end users would have been seriously impacted**

A Production Use Case – XenApp Server Abnormal Behavior

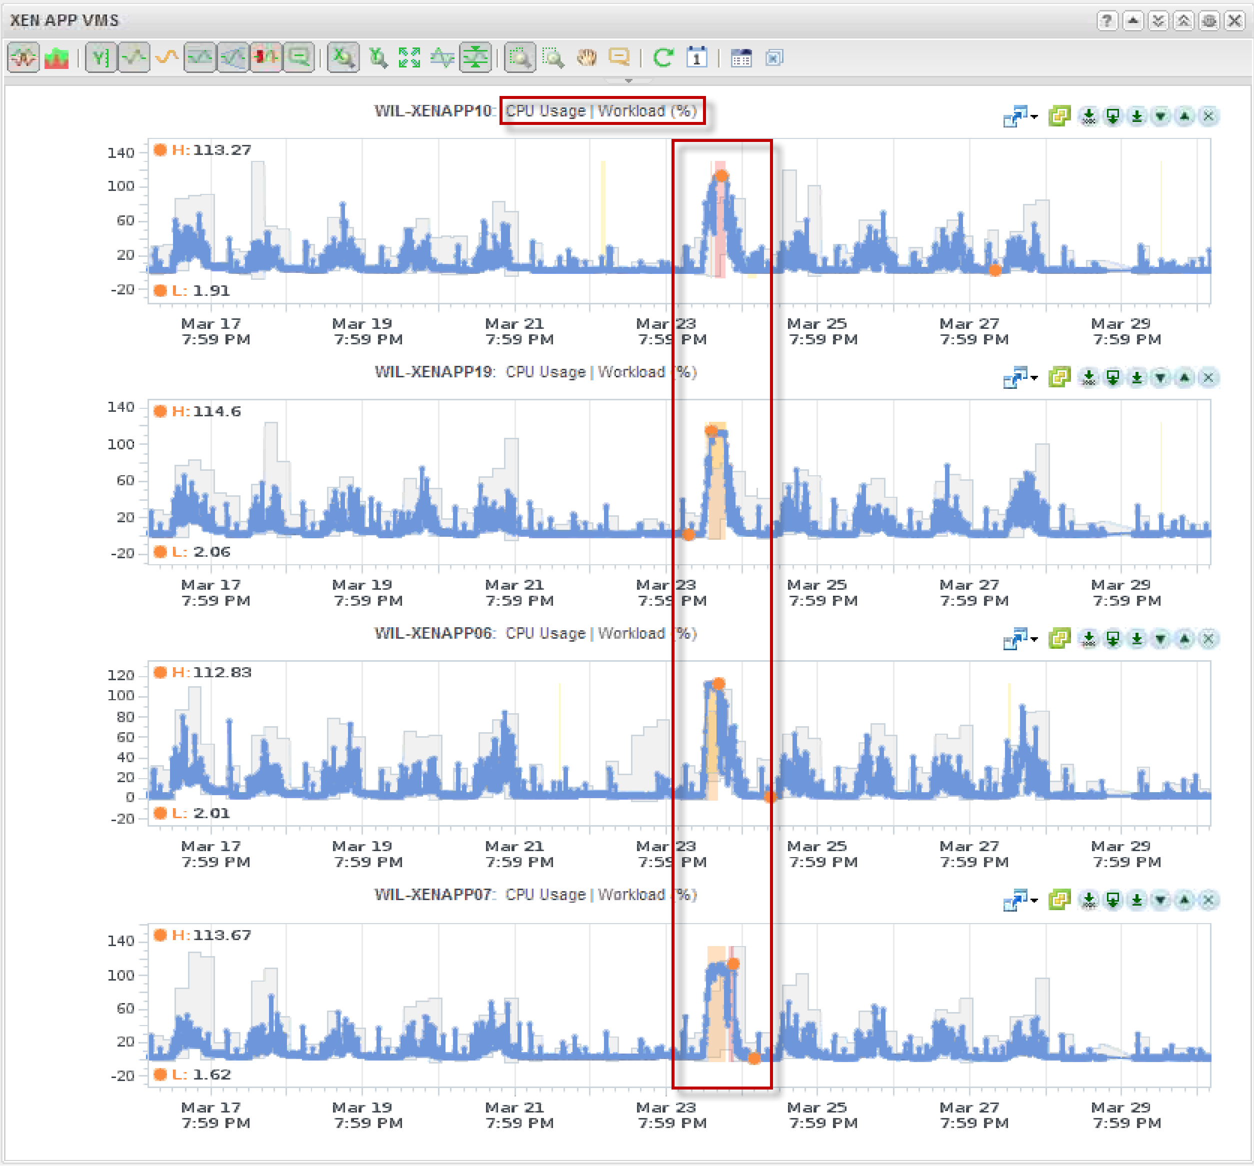

vmware

A Production Use Case – XenApp Server Abnormal Behavior

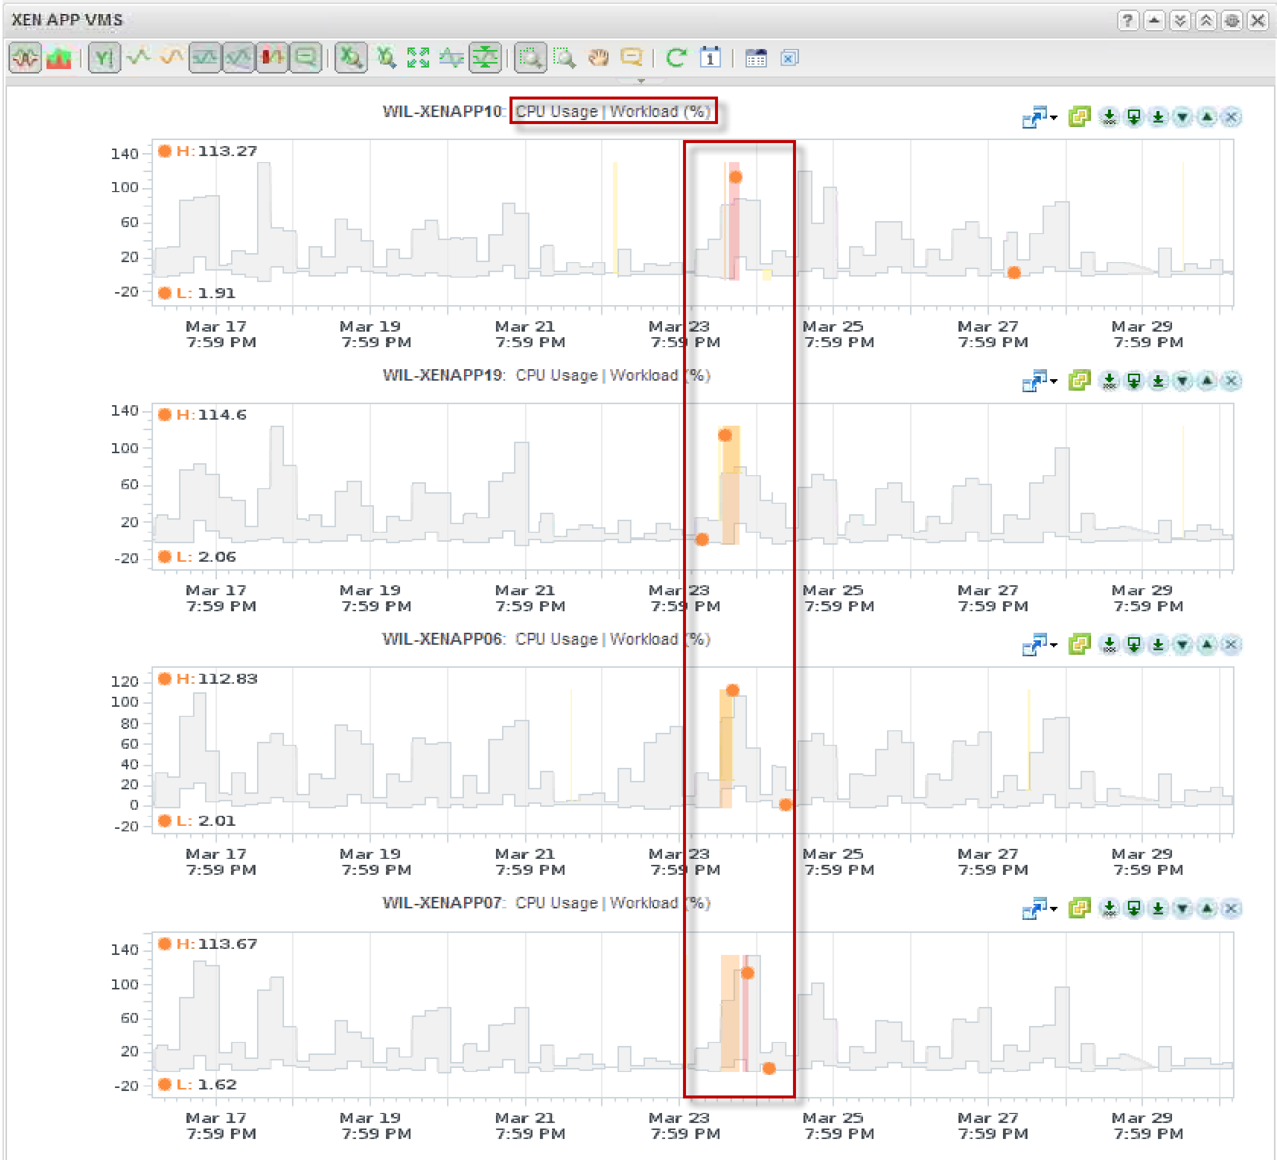

Conclusions

- Our categorization techniques allow achieving a more accurate Dynamic Threshold for the individual metric
- By using optimization techniques we achieve optimal balance between false positive and false negative alerts
- This would not be possible with classical parametric approaches including Fourier transform, and other common purpose enterprise algorithms
- Moreover, this approach enables other algorithms to be applied to determine system issues with more accuracy.