

Predictive Analysis for Migration Schedulers

Steven Rosenberg

Virtualization and SLA R&D Department

Red Hat

Overview:

Load Balancing

Types of Solutions

Fault Tolerance

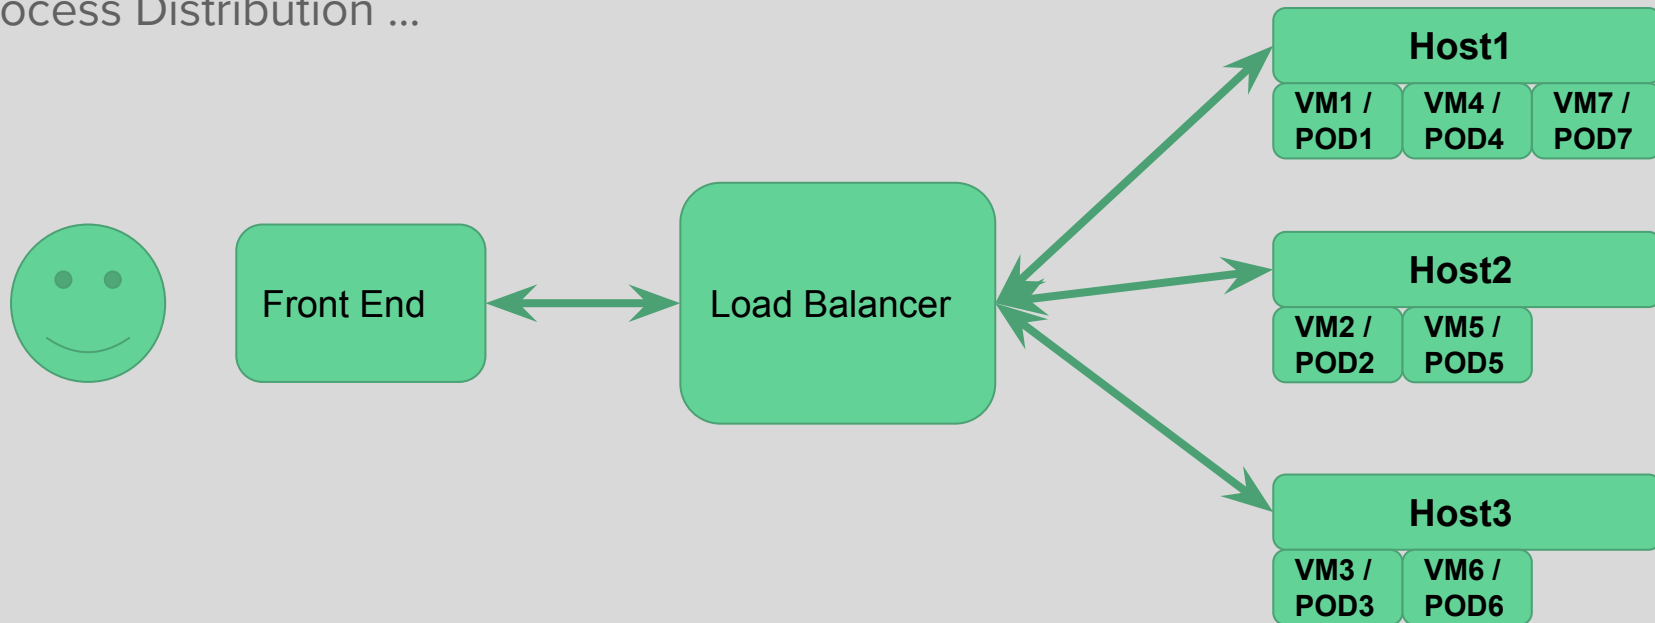
Live Migration

Scheduling

Predictive Analysis

Load Balancing Example

Process Distribution ...



Load Balancing


- Priority Based Upon Urgency
- Even Distribution within categories:
 - Urgent Priority – Mission Critical – Real Time Processing
 - High Priority – High Importance – near Real Time Processing
 - Neutral Priority – Medium Importance – Normal Processing
 - Low Priority – Low Importance - Not Time Critical Processing
 - No Priority – Unimportant – Unimportant Processes

Fault Tolerance

- Failure Types:
 - Network Element Failures
 - Hardware / Resource Failures
 - OS / BIOS / Kernel Failures
 - Process Failures

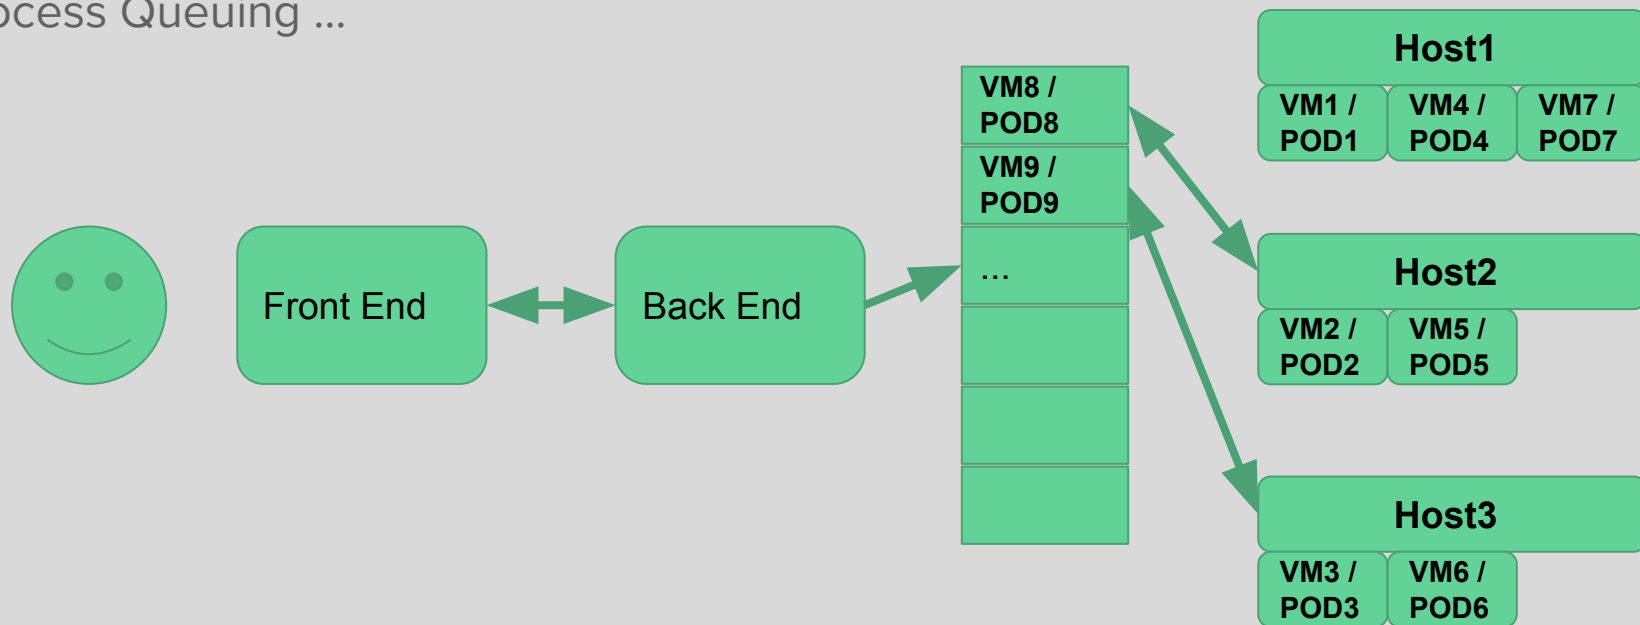
Fault Tolerance Redundancy Example

High Availability ...



Scheduler Dispatching Concepts

Process Queuing ...



Scheduling

- Ability to launch processes based upon needed resources
 - Monitor the amount of resources each process utilizes
- Types of Launching/Migration Scenarios:
 - Initial Launch
 - Migration for Maintenance
 - Re-balancing - Migration to Another Host
 - Fault Recovery – Migrating to mitigate system/process failure.

Policy Units - Attributes of Scheduling Migrations

- Filters
- Weights/Scoring
- Balancers:
 - Even Distribution
 - Power Saving
 - Prioritizing
 - Affinity
 - CPU/NUMA Pinning for Optimal Performance

Types of Solutions For Applying Predictive Analysis:

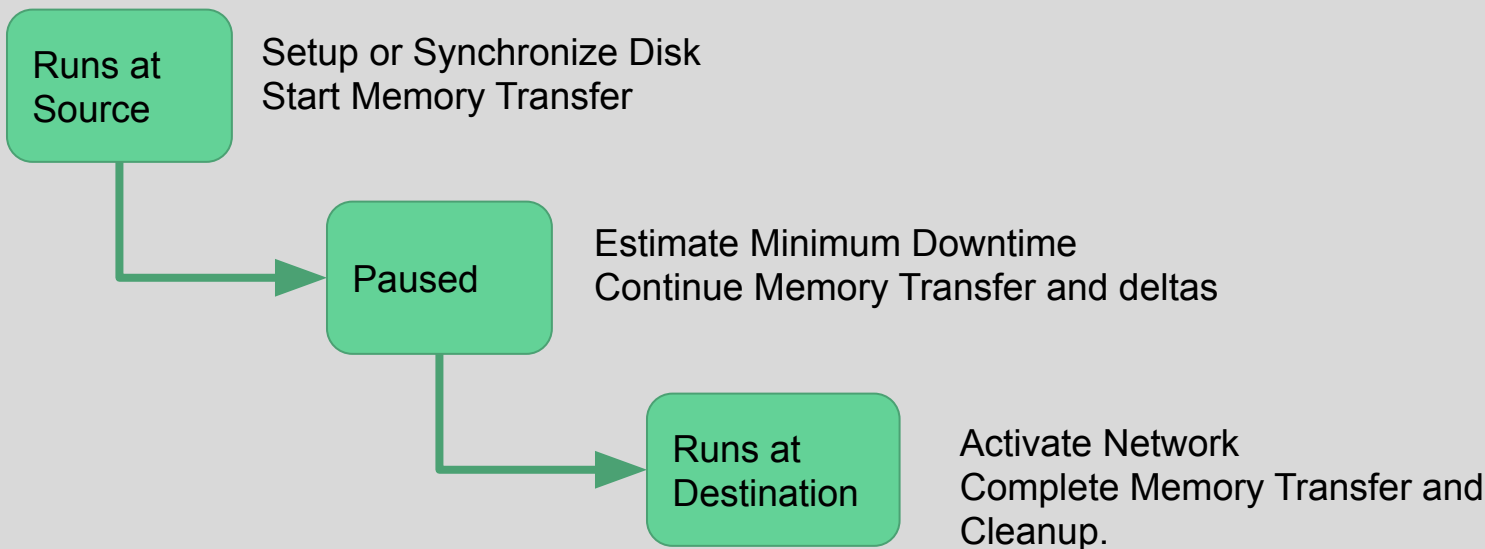
- Live Migration
 - Load Balancing
 - Fault Recovery
 - Minimizing Live Migration Pausing
- Redundancy
 - Distribution of processes running simultaneously
 - Fault Recovery

Live Migration:

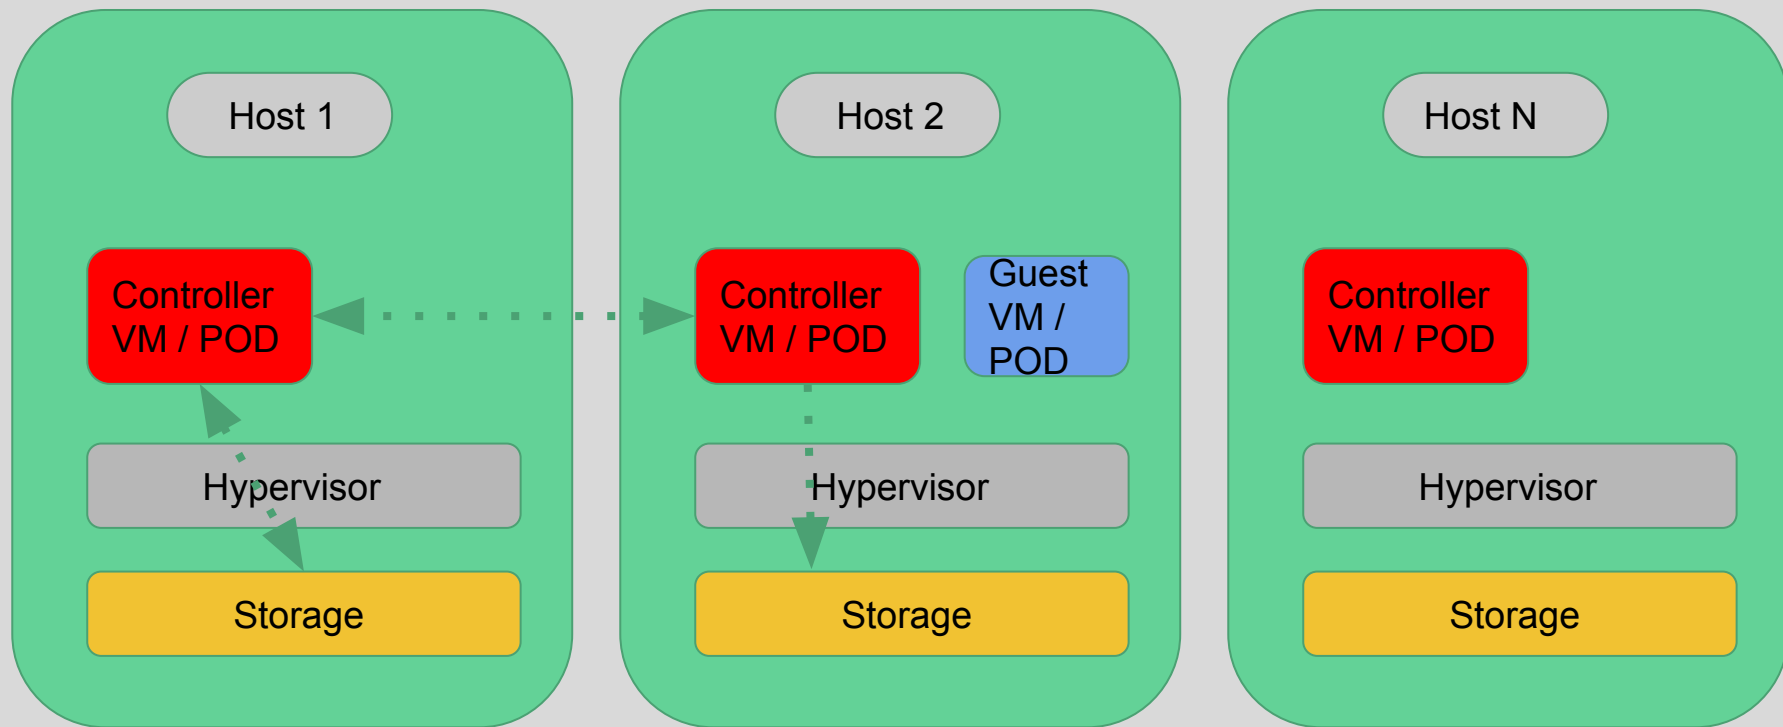
- Network Connectivity.
- Remote Disk(s) Availability
- Migrating Data on Local Disk(s)
- Copying Memory State in phases
 - All of the current memory contents
 - Current Differences before VM / POD Pausing
 - Minimal Differences during VM / POD Pausing
- Copy CPU State
- The goal is to limit pausing of the VM / POD
- Restarting the VM /POD on the Destination Host
- Cleanup on the Source Host

Live Migration Transitioning

Sequence of Events ...



Live Migration From Host 1 to 2 Transitioning




Predictive Analysis Topics For Discussion

- Predicting future occurrences via analysis of past performance
- Techniques for Predictive Analysis
- Process for Developing a Prediction Model
- Types of Predictive Models with Examples
- Applying to Scheduling

Predictive Analytics Methodology

Modeling ...



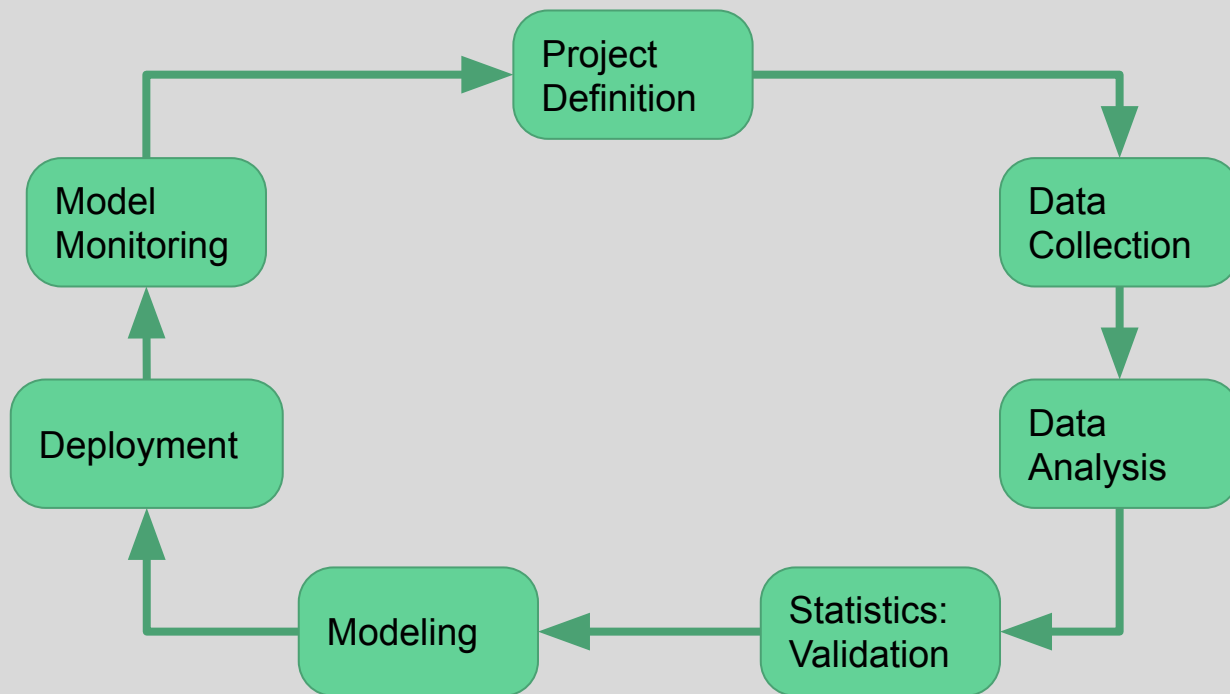
Techniques for Predictive Analysis

- Regression techniques
- Linear regression model
- Discrete choice models
- Logistic regression
- Multinomial logistic regression
- Probit regression
- Logit versus probit
- Time series models
- Survival or duration analysis
- k-nearest neighbours

Classification and regression trees (CART)
Multivariate adaptive regression splines
Multivariate adaptive regression splines
Machine learning techniques
Neural networks
Multilayer perceptron (MLP)
Radial basis functions
Support vector machines
Naïve Bayes
Geospatial predictive modeling

Process for Developing a Prediction Model

Cyclical ...



Types of Predictive Models with Examples

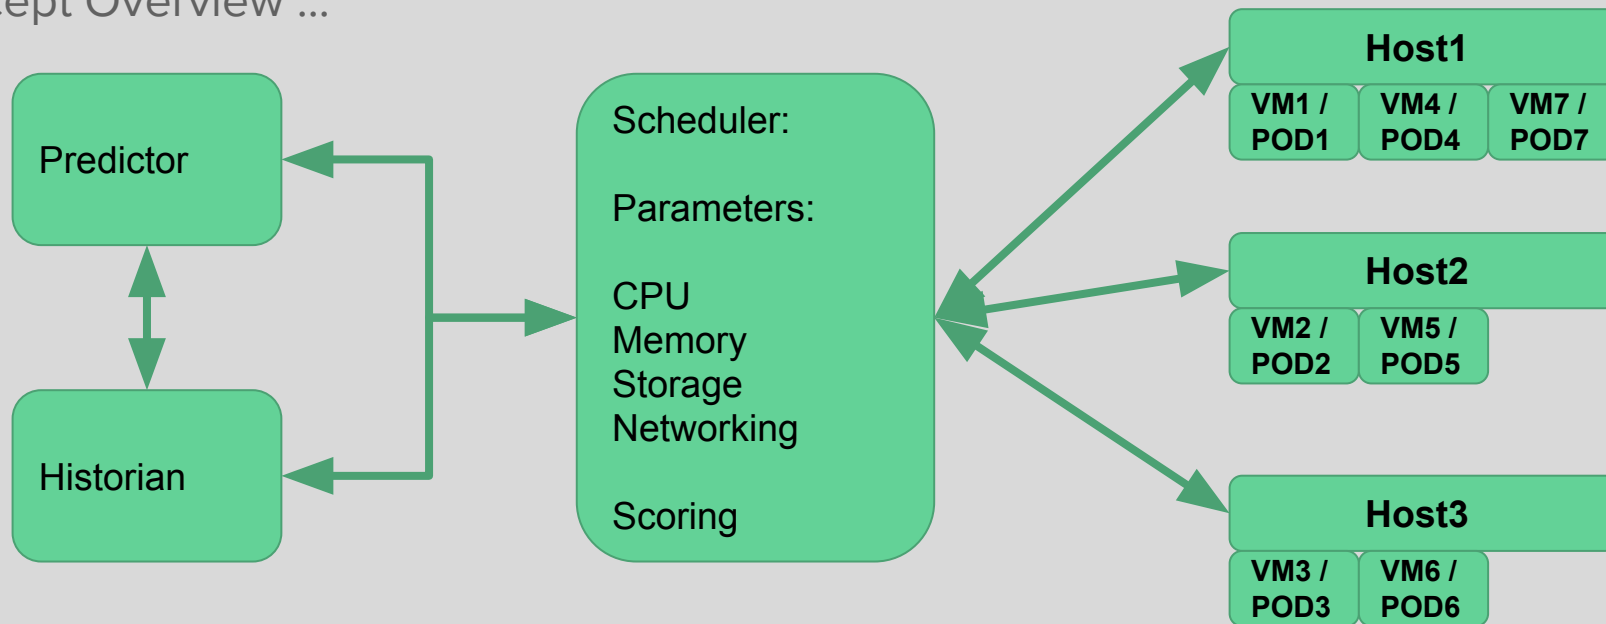
- Support Vector Machine – Model
 - Classification – To predict a category
 - Example: Stock prices increase or decrease – Yes or No, True or False answer.
- Predict Quantity – Regression
 - Example: Predicting a person's age based upon height, weight, health and other factors
- Detect an Anomaly – Normal Behavior or Exception (Anomaly)
 - Example: Money withdrawal anomalies
- Clustering: Discover Structure in Unexplored Data
 - Example: Finding groups of customers with similar behavior given a large Database of customers containing their demographics and past buying records.

Applying Predictive Analytics to Schedulers

- Criteria for Data
 - Processing Time / Iterations - Adjusted for Resource Capacity and Priority
 - % of Resources used – Adjusted for Capacity and Priority
 - Adjust for anomalies when Calculating Averages
- Ideas - Selective Techniques applied for other scheduling applications:
 - Combining regression-like modeling and functional approximation, using the sum of exponential functions, to produce probability estimates.
 - Machine Learning & Advanced Mathematical Models.

Predictive Analysis Architecture

Concept Overview ...



Tracking Historical Data

- The Time Each Process Starts and Terminates
- The Resources Used By Each Process
- The Time Each Process Uses To Migrate
- The Time / Iterations that Memory / Disk Transfer Occurs Per Size

Considerations Based Upon Analysis

- If Early Migration Can Proceed
- When Early Migration Shall Start
- Error Correction / Anomaly Detection for Accurate Results

Anomaly / Error Calculation Methods to Consider

- Statistical - Calculate % of Error From the Mean and Eliminate Results Outside of the Threshold
- Signal Processing Techniques - Smoothing Filter to Eliminate Glitches
- Machine Learning - Analysis of Patterns and Categorize Between Normal And Out Of Range Results.

Thank You

Steven Rosenberg

srosenbe@redhat.com