

Appendix

ACIT5900

Master Thesis

in

Applied Computer and Information

Technology (ACIT)

May 2023

Cloud-Based Services and Operations

Towards Green Container Management:

A novel approach for container resource
allocation through statistical modeling

Håkon Borgersen Ay

Department of Computer Science

Faculty of Technology, Art and Design

OSLOMET

© 2023 Håkon Borgersen Ay - s360599

Towards Green Container Management: A novel approach for container resource allocation through statistical modeling

Supervisor: Kyrre Begnum

<https://www.oslomet.no>

Printed: Oslo Metropolitan University

Contents

Contents	ii
Figures	iii
Code Listings	iv
A Survey Sent to Developers at Intility	1
B Python Code	7
B.1 Pre-processing Code	7
B.1.1 Reading and pre-processing .csv files	7
B.1.2 Converting memory to float	8
B.1.3 Calculating skewness	9
B.1.4 Script for showing daily memory usage during a 2 week period	10
B.2 Linear programming	11
B.2.1 Calculating The MREI	13
B.3 ARIMA	13
B.3.1 ACF and PACF	14
B.3.2 Fitting the ARIMA model	14
B.3.3 Test Data Underprediction Rate (TDUR)	15
B.3.4 MAE, MAPE and RMSE	16
B.4 Facebook's Prophet Model	16
B.5 P+LP	17
B.6 Performance Score	19

Figures

A.1	Survey Sent to Developers at Intility: Question 1	2
A.2	Survey Sent to Developers at Intility: Question 2	2
A.3	Survey Sent to Developers at Intility: Question 3	3
A.4	Survey Sent to Developers at Intility: Question 4	3
A.5	Survey Sent to Developers at Intility: Question 5	4
A.6	Survey Sent to Developers at Intility: Question 6	4
A.7	Survey Sent to Developers at Intility: Question 7	5
A.8	Survey Sent to Developers at Intility: Question 8	5
A.9	Survey Sent to Developers at Intility: Question 9	6
A.10	Survey Sent to Developers at Intility: Question 10	6

Code Listings

B.1	Reading and appending .csv files	7
B.2	Converting memory to float	8
B.3	Calculating skewness in container	9
B.4	Script showing daily memory usage for a 2 week period	10
B.5	The Linear Programming Model	11
B.6	Calculating the MREI	13
B.7	Determining differencing for the ARIMA model	13
B.8	Calculating ACF and PACF for ARIMA	14
B.9	Fitting the Arima model	14
B.10	Calculating TDUR	15
B.11	Calculating MAE, MAPE, and RMSE	16
B.12	Facebook's Prophet Model	16
B.13	Finalized P+LP model, included plotting	17
B.14	Calculating Performance Score	19

Appendix A

Survey Sent to Developers at Intility

Appendix B consists of all questions and answers from the survey sent out to developers at Intility. Below is the description of the survey which was sent out together with the survey

Hello, my name is Håkon Borgersen Ay. I work part-time at the developer infrastructure team and will be starting full-time as a trainee in August 2023. Right now I'm writing my master thesis with the working title: "Optimizing Resource Utilization in a Kubernetes Cluster: Investigating and Automating Resource Allocation using Workload Profiles".

I am conducting an anonymous survey to gather information about resource utilization and cloud waste associated with cloud resources. Your input will help me identify areas of improvement and potential cost savings.

The term Cloud is in this case used to describe Kubernetes clusters. Consider the resource usage vs. the allocated resources of clusters, nodes, and pods when answering. When the word resources are used I'm referring to the allocated CPU and RAM.

Please note that while some of the data collected through this survey may be used in my master thesis, certain information may be kept confidential due to policy constraints. I appreciate your participation.

Do you expect your spend in Kubernetes clusters to increase, decrease, or stay the same over the next 12 months?

9 svar




Figure A.1: Survey Sent to Developers at Intility: Question 1

What percentage of resources allocated in the Kubernetes cluster do you believe is wasted, i.e. spent on unused or idle resources? (e.g. x%)

7 svar



Figure A.2: Survey Sent to Developers at Intility: Question 2

How confident are you that your organization knows with certainty how much of the clusters' resources is wasted?

8 svar




Figure A.3: Survey Sent to Developers at Intility: Question 3

What would you say is the impact of your organization's resource waste? Select all that apply.

8 svar




Figure A.4: Survey Sent to Developers at Intility: Question 4

Answers from Figure A.4

- Reduces our profitability
- Makes IT look bad to the rest of the organization
- Makes us less competitive
- Difficulty staying within budget
- Can't hire resources we need due to budget constraints
- Bad for the environment
- Clutter, unused capacity

Is reducing resource waste a high priority for your organization?
9 svar




Figure A.5: Survey Sent to Developers at Intility: Question 5

What are the biggest causes of resource waste for your organization, in your opinion? Select all that apply.
9 svar




Figure A.6: Survey Sent to Developers at Intility: Question 6

Answers from figure A.6:

- Cloud complexity makes it hard to estimate how many resources are actually needed
- Over-provisioning to ensure applications perform well
- Maintaining more environments than needed
- Developers have no incentive to run apps efficiently
- Not enough attention or management oversight of cloud spend
- No visibility into actual resource utilization vs. what we're paying for
- Controlling cloud spend is not a priority for us
- Ensuring we are ready for a potential surge in demand
- Applications are configured for high availability, causing us to run more resources than actually needed

If you are currently using Kubernetes, do you believe the complexity of Kubernetes contributes to your organization's cloud waste?

8 svar




Figure A.7: Survey Sent to Developers at Intility: Question 7

How are Kubernetes resource allocation decisions generally made in your organization?

9 svar




Figure A.8: Survey Sent to Developers at Intility: Question 8

Answers from Figure A.8

Relying on defaults from vendors or service providers

Optimization based on machine learning

Best guess / Trial and error

Don't know

Kubectl top pods and guesses

Using statistics from previously running workloads, while over-provisioning to ensure application performance

Best practices from vendor for 3rd party applications, best guess for in-house developed applications

How has the resource allocation decisions impacted and affected the performance and reliability of running services

9 svar




Figure A.9: Survey Sent to Developers at Intility: Question 9

Are there any specific features or capabilities you would like to see added to your current cloud resources?

4 svar

- Comprehensive monitoring tool would be nice
- Kubernetes Vertical Pod Autoscaler
- Optimization based on machine learning or more defaults/best practice options from intility
- Easier auto-scaling options.

Figure A.10: Survey Sent to Developers at Intility: Question 10

Appendix B

Python Code

B.1 Pre-processing Code

B.1.1 Reading and pre-processing .csv files

Code listing B.1: Reading and appending .csv files

```
1 # Reading and appending .csv files (ram)
2
3 container=1
4 df1 = pd.read_csv(f'B{container}.csv',
5                 sep=',',
6                 low_memory=False)
7
8 orig_name= df1.columns[1]
9 df1 = df1.rename(columns={orig_name: 'Containers: Memory usage'})
10
11 # Find the index of the last non-null value
12 last_non_null_index = df1['Containers: Memory usage'].last_valid_index()
13
14 # Remove rows with missing values after the last non-null value
15 df1 = df1.loc[:last_non_null_index]
16
17 # Counting missing values and find the percentage
18 #df1.info()
19 #print(df1['Containers: Memory usage'].value_counts())
20 #print(df1[' bytes - collector-xcfsq collector | collector'].value_counts())
21
22 # Standardized date format
23 df1['Date'] = pd.to_datetime(df1['Date'])
24
```

```

25 # converting n/a's to mean value
26 df1['Containers: Memory usage'] = df1['Containers: Memory usage'].apply(
    convert_memory_to_float)
27 df1['Containers: Memory usage'] = df1['Containers: Memory usage'].fillna(df1[
    'Containers: Memory usage'].mean())
28
29 fig, ax = plt.subplots(figsize=(12, 6))
30 ax.set_xlabel('Date')
31 ax.set_ylabel('Memory Usage (MB)')
32 ax.set_title('Memory Usage Over Time container01')
33 ax.grid()
34
35 # Customize date ticks and format
36 ax.xaxis.set_major_locator(mdates.AutoDateLocator())
37 ax.xaxis.set_major_formatter(mdates.DateFormatter('%Y-%m-%d'))
38 plt.setp(ax.get_xticklabels(), rotation=30, ha='right')
39
40 #plt.savefig(f'memory_usage_c{container}.png', dpi=300, bbox_inches='tight')
41 plt.show()

```

B.1.2 Converting memory to float

Code listing B.2: Converting memory to float

```

1 def convert_memory_to_float(value):
2     # Check if the value is already numeric
3     if isinstance(value, (int, float)):
4         return float(value)
5
6     # If value is a string, check the unit
7     if isinstance(value, str):
8         unit = None
9         number = value.strip()
10        if 'B' in value.upper() and 'MB' not in value.upper() and 'KB' not in value
        .upper():
11            unit = 'B'
12            number = value[:-1].strip()
13        elif 'KB' in value.upper():
14            unit = 'KB'
15            number = value[:-2].strip()
16        elif 'MB' in value.upper():
17            unit = 'MB'
18            number = value[:-2].strip()
19        elif 'GB' in value.upper():

```

```

20         unit = 'GB'
21         number = value[:-2].strip()
22
23         # Convert the numeric part to a float
24         value = float(number)
25
26         # Convert the value to MB based on the unit
27         if unit == 'B':
28             value = value / (1024 * 1024)
29         elif unit == 'KB':
30             value = value / 1024
31         elif unit == 'GB':
32             value = value * 1024
33
34     return value

```

B.1.3 Calculating skewness

Code listing B.3: Calculating skewness in container

```

1 column_name = "Container: Memory usage"
2
3 skewness = df1['Containers: Memory usage'].skew()
4
5 print(df1['Containers: Memory usage'].median())
6 # Create a histogram with 10 bins and a density plot
7 fig, ax = plt.subplots()
8 ax.hist(df1['Containers: Memory usage'], bins=20, density=True, alpha=0.5, color="
    blue")
9 ax.set_xlabel(column_name)
10 ax.set_ylabel("Density")
11
12 ax.set_xlim(100, max(df1['Containers: Memory usage']))
13
14 # Add a vertical line at the median, mean, mode, and skewness value
15 ax.axvline(df1['Containers: Memory usage'].median(), color="red", linestyle="--",
    label="Median")
16 ax.axvline(df1['Containers: Memory usage'].mean(), color="orange", linestyle="--",
    label="Mean")
17 ax.axvline(df1['Containers: Memory usage'].mode()[0], color="green", linestyle="--"
    , label="Mode")
18 ax.axvline(skewness, color="purple", linestyle="--", label="Skewness: {:.2f}".
    format(skewness))
19

```

```

20 # Add a legend and title
21 ax.legend()
22 ax.set_title("Distribution of memory usage")
23
24 # Display the plot
25 #plt.savefig(f'memory_usage_distribution_c{container}.png', dpi=300, bbox_inches='
    tight')
26 plt.show()

```

B.1.4 Script for showing daily memory usage during a 2 week period

Code listing B.4: Script showing daily memory usage for a 2 week period

```

1
2 def plot_24_hour_segments(df, column):
3     days = df.index.normalize().unique()
4
5     plt.figure(figsize=(12, 6))
6
7     for day in days:
8         day_data = df.loc[day:day + pd.Timedelta('1D') - pd.Timedelta('5m'), column
9             ]
10        plt.plot(day_data.index.time, day_data.values, label=day.date())
11
12 df_days=df1
13 df_days = df_days[['Date', 'Containers: Memory usage']]
14 df_days['Date'] = pd.to_datetime(df_days['Date'])
15 df_days.set_index('Date', inplace=True)
16
17
18
19 def plot_24_hour_segments(df, column, window_size=12):
20     first_timestamp = df.index[0]
21     days = df.index.normalize().unique()
22
23     # Calculate end date
24     last_full_day = df.index[-1].normalize() - pd.Timedelta('1D')
25     end_date = last_full_day + pd.Timedelta(hours=first_timestamp.hour, minutes=
        first_timestamp.minute)
26
27     # Apply rolling average
28     df_smooth = df.rolling(window=window_size, center=True).mean()
29

```

```

30 plt.figure(figsize=(12, 6))
31
32 for day in days:
33     start_time = day + pd.Timedelta(hours=first_timestamp.hour, minutes=
34         first_timestamp.minute)
35     end_time = start_time + pd.Timedelta('1D') - pd.Timedelta('5m')
36
37     if end_time <= end_date:
38         day_data = df_smooth.loc[start_time:end_time, column]
39
40         if not day_data.empty:
41             times = [t.hour * 60 + t.minute for t in day_data.index.time]
42             plt.plot(times, day_data.values, label=day.date())
43
44 plt.xlabel('Time of Day')
45 plt.ylabel('Memory Usage in MiB')
46 plt.legend()
47 plt.title('Daily Memory Usage Patterns (Smoothed)')
48
49 one_day_minutes = 24 * 60
50 plt.xlim(0, one_day_minutes)
51
52 # Set x-axis ticks and labels
53 x_ticks = np.arange(0, one_day_minutes + 1, 2 * 60) # Every 4th hour in
54     minutes
55 x_labels = [f'{t // 60:02d}:00' for t in x_ticks]
56 plt.xticks(x_ticks, x_labels, rotation=45)
57
58 plt.savefig('daily_memory_usage_patterns.png', dpi=300, bbox_inches='tight')
59 plt.show()
60 plot_24_hour_segments(df_days, 'Containers: Memory usage')

```

B.2 Linear programming

Code listing B.5: The Linear Programming Model

```

1
2 # Linear programming modeling
3 import pulp
4
5 safety_margin_percentage = 10 # Additional memory allocated as a percentage of the
6     max usage

```



```
6 safety_margin = max(df1['Containers: Memory usage']) * (safety_margin_percentage /
7     100)
8 # Calculate the typical memory usage (e.g., the mean or median)
9 typical_memory_usage = np.mean(df1['Containers: Memory usage']) * 1.05
10
11 # Defining the problem, and specifying that the objective is to minimize a function
12 prob = pulp.LpProblem("ContainerResourceAllocation", pulp.LpMinimize)
13
14 # Variables
15 x_request = pulp.LpVariable("x_request", lowBound=typical_memory_usage)
16 x_limit = pulp.LpVariable("x_limit", lowBound=max(df1['Containers: Memory usage'])
17     + safety_margin)
18
19 # Objective function
20 prob += x_request
21
22 # Constraints
23 prob += x_request <= x_limit
24
25 # Solve the problem
26 status = prob.solve()
27
28 # Print the results
29 #print("Status:", pulp.LpStatus[status])
30 req_value = pulp.value(x_request)
31 lim_value = pulp.value(x_limit)
32 print("Request limit for Container 1:", pulp.value(x_request))
33 print("Memory limit for Container 1:", pulp.value(x_limit))
34
35 fig, ax = plt.subplots(figsize=(12, 6))
36 ax.plot(df1['Date'], df1['Containers: Memory usage'])
37 ax.set_xlabel('Date')
38 ax.set_ylabel('Memory Usage (MB)')
39 ax.set_title(f'Memory Usage Over Time container0{container}')
40 ax.grid()
41
42 # Customizing date ticks and format
43 ax.xaxis.set_major_locator(mdates.AutoDateLocator())
44 ax.xaxis.set_major_formatter(mdates.DateFormatter('%Y-%m-%d'))
45 plt.setp(ax.get_xticklabels(), rotation=30, ha='right')
46
47 ax.axhline(y=req_value, color='r', linestyle='--', label='Request')
48 ax.axhline(y=lim_value, color='b', linestyle='--', label='Limit')
```

```
49
50 ax.legend()
51
52 # print plot
53 plt.savefig(f'memory_usage_linear_programming_c{container}', dpi=300, bbox_inches='
    tight')
54 # Show the plot
55 plt.show()
```

B.2.1 Calculating The MREI

Code listing B.6: Calculating the MREI

```
1 def calculate_kpi(df, request_value):
2     """
3     Calculate the percentage of container usage memory datapoints under the set
4         request.
5
6     :param df: DataFrame with container memory usage data
7     :param request_value: Request limit for the container
8     :return: KPI as a percentage
9     """
10    under_request = df[df['y'] < request_value]
11    percentage = (len(under_request) / len(df)) * 100
12    return percentage
```

B.3 ARIMA

Determining Differencing

Code listing B.7: Determining differencing for the ARIMA model

```
1
2 adf_p_value_0 = sm.tsa.stattools.adfuller(df1['Containers: Memory usage'])[1]
3 adf_p_value_1 = sm.tsa.stattools.adfuller(df1['Containers: Memory usage'].diff().
4     dropna())[1]
5 adf_p_value_2 = sm.tsa.stattools.adfuller(df1['Containers: Memory usage'].diff().
6     diff().dropna())[1]
7
8 # Create a single figure with the desired layout
9 fig, axes = plt.subplots(nrows=2, ncols=3, figsize=(18, 8))
10
11 # Without Differencing
```

```

10 axes[0, 0].set_title(f'Without Differencing (p={adf_p_value_0:.2e})')
11 axes[0, 0].plot(df1['Containers: Memory usage'])
12 sm.graphics.tsa.plot_acf(df1['Containers: Memory usage'].dropna(), ax=axes[1, 0])
13
14 # 1st Order Differencing
15 axes[0, 1].set_title(f'1st Order Differencing (p={adf_p_value_1:.2e})')
16 axes[0, 1].plot(df1['Containers: Memory usage'].diff())
17 sm.graphics.tsa.plot_acf(df1['Containers: Memory usage'].diff().dropna(), ax=axes
    [1, 1])
18
19 # 2nd Order Differencing
20 axes[0, 2].set_title(f'2nd Order Differencing (p={adf_p_value_2:.2e})')
21 axes[0, 2].plot(df1['Containers: Memory usage'].diff().diff())
22 sm.graphics.tsa.plot_acf(df1['Containers: Memory usage'].diff().diff().dropna(), ax
    =axes[1, 2])
23
24 plt.savefig(f'memory_usage_plots_with_adf_c{container}.png', dpi=300, bbox_inches='
    tight')
25
26 # Show the plots
27 plt.show()

```

B.3.1 ACF and PACF

Code listing B.8: Calculating ACF and PACF for ARIMA

```

1 stationary_data = df1['Containers: Memory usage'].diff().dropna()
2 # Create ACF and PACF plots
3 fig, axes = plt.subplots(2, 1, figsize=(12, 8))
4
5 sm.graphics.tsa.plot_acf(stationary_data, lags=40, ax=axes[0])
6 sm.graphics.tsa.plot_pacf(stationary_data, lags=40, ax=axes[1])
7
8 plt.savefig(f'memory_usage_plots_PACFnACF_c{container}.png', dpi=300, bbox_inches='
    tight')
9 plt.show()

```

B.3.2 Fitting the ARIMA model

Code listing B.9: Fitting the Arima model

```

1
2 memory_usage = df1['Containers: Memory usage']

```

```

3
4 train_size = int(len(memory_usage) * 0.80) # 80% of the data for training
5 train_data = memory_usage[:train_size]
6 test_data = memory_usage[train_size:]
7
8 p,d,q = 1,2,1
9
10 model = ARIMA(train_data, order=(p, d, q))
11 results = model.fit()
12
13 predictions = results.predict(start=train_size, end=len(memory_usage) - 1, dynamic=
    True)
14
15 # Plot the training data, test data, and predictions
16 fig, ax = plt.subplots(figsize=(12, 6))
17 ax.set_xlabel('Date')
18 ax.set_ylabel('Memory Usage (MB)')
19 ax.set_title(f'Memory Usage Over Time with ARIMA Predictions for c{container}')
20 ax.grid()
21
22 # Customizing date ticks and format
23 ax.xaxis.set_major_locator(mdates.AutoDateLocator())
24 ax.xaxis.set_major_formatter(mdates.DateFormatter('%Y-%m-%d'))
25 plt.setp(ax.get_xticklabels(), rotation=30, ha='right')
26
27 ax.plot(df1['Date'][:train_size], train_data, label='Training Data')
28 ax.plot(df1['Date'][train_size:], test_data, label='Test Data')
29 ax.plot(df1['Date'][train_size:], predictions, label='Predictions')
30
31 ax.legend()
32
33 # Save and show the plot
34 plt.savefig(f'arima_memory_usage_c{container}.png', dpi=300, bbox_inches='tight')
35 plt.show()

```

B.3.3 Test Data Underprediction Rate (TDUR)

Code listing B.10: Calculating TDUR

```

1
2 def calculate_kpi_predictions(df, predictions):
3     """
4     Calculate the number of container usage memory datapoints under the predicted
        values.

```

```

5
6     :param df: DataFrame with container memory usage data
7     :param predictions: Predicted values from the ARIMA model
8     :return: Number of datapoints under the predicted line
9     """
10    test_data = df[-len(predictions):] # Get the test data corresponding to the
        predictions
11    under_prediction = test_data[test_data['Containers: Memory usage'] <
        predictions]
12    percentage = (len(under_prediction) / len(test_data)) * 100
13    return percentage
14
15 # Assuming the 'predictions' variable contains the ARIMA model predictions
16 kpi = calculate_kpi_predictions(df1, predictions)
17 print(kpi)

```

B.3.4 MAE, MAPE and RMSE

Code listing B.11: Calculating MAE, MAPE, and RMSE

```

1
2 def mean_absolute_percentage_error(y_true, y_pred):
3     y_true, y_pred = np.array(y_true), np.array(y_pred)
4     return np.mean(np.abs((y_true - y_pred) / y_true)) * 100
5
6 mae = np.mean(np.abs(predictions - test_data))
7 mape = mean_absolute_percentage_error(test_data, predictions)
8 rmse = np.sqrt(np.mean((test_data - predictions) ** 2))

```

B.4 Facebook's Prophet Model

Code listing B.12: Facebook's Prophet Model

```

1 model = Prophet()
2
3 df = pd.DataFrame({'ds': df1['Date'], 'y': df1['Containers: Memory usage']})
4 df['ds'] = pd.to_datetime(df['ds'])
5
6 # Split the data into train and test sets
7 train_size = int(len(df) * 0.8) # 80% of the data for training, 20% for testing
8 train_df = df[:train_size]
9 test_df = df[train_size:]

```

```

10
11
12 # Fit the model on the training data
13 model.fit(train_df)
14
15 # Make predictions on the test data
16 future = model.make_future_dataframe( periods=len(test_df), freq='5min',
17                                     include_history=False)
18 forecast = model.predict(future)
19
20 # Plot the forecast
21 fig = model.plot(forecast)
22 ax = fig.gca()
23 ax.set_title(f'Prophet Forecast for c{container}')
24 ax.set_xlabel('Date')
25 ax.set_ylabel('Memory Usage (MB)')
26
27 # Plot the test data
28 ax.scatter(test_df['ds'], test_df['y'], color='purple', label='Test Data',s=10)
29
30 # Add a legend
31 ax.legend(['Observed', 'Forecast', 'Uncertainty Interval', 'Test Data'])
32 plt.savefig(f'prophetC{container}.png', dpi=300, bbox_inches='tight')
33
34 plt.show()

```

B.5 P+LP

Code listing B.13: Finalized P+LP model, included plotting

```

1 model = Prophet(changepoint_prior_scale=0.01)
2
3 df = pd.DataFrame({'ds': df1['Date'], 'y': df1['Containers: Memory usage']})
4 df['ds'] = pd.to_datetime(df['ds'])
5
6 orig_req=15
7 orig_lim=15
8
9 # Split the data into train and test sets
10 train_size = int(len(df) * 0.8) # 80% of the data for training, 20% for testing
11 train_df = df[:train_size]
12 test_df = df[train_size:]

```

```
13
14 # Fit the model on the training data
15 model.fit(train_df)
16
17 # Make predictions on the test data
18 future = model.make_future_dataframe(periods=len(test_df), freq='60min',
    include_history=False)
19 forecast = model.predict(future)
20
21 # Plot the forecast
22 fig = model.plot(forecast)
23 ax = fig.gca()
24 ax.set_title(f'P+LP Forecast for c{container}')
25 ax.set_xlabel('Date')
26 ax.set_ylabel('Memory Usage (MB)')
27
28 # Plot the test data
29 ax.scatter(test_df['ds'], test_df['y'], color='purple', label='Test Data', s=10)
30
31 # Calculate the new request and limit lines
32 new_request = np.mean(forecast['yhat']) * 1.05 # Adding 10% safety margin
33
34 max_observed = max(train_df['y'])
35 max_predicted = max(forecast['yhat_upper'])
36 higher_value = max(max_observed, max_predicted)
37 new_limit = higher_value * 1.15 # 10% over the higher value
38
39 # Old request and limit lines calculation
40 safety_margin_percentage = 15 # Additional memory allocated as a percentage of the
    max usage
41 safety_margin = max(train_df['y']) * (safety_margin_percentage / 100)
42 typical_memory_usage = np.mean(train_df['y']) * 1.05
43
44 # Define the problem
45 prob = pulp.LpProblem("ContainerResourceAllocation", pulp.LpMinimize)
46
47 # Variables
48 x_request = pulp.LpVariable("x_request", lowBound=typical_memory_usage)
49 x_limit = pulp.LpVariable("x_limit", lowBound=max(train_df['y']) + safety_margin)
50
51 # Objective function
52 prob += x_request
53
54 # Constraints
55 prob += x_request <= x_limit
```

```
56
57 # Solve the problem
58 status = prob.solve()
59
60 # Get the old request and limit values
61 req_value = pulp.value(x_request)
62 lim_value = pulp.value(x_limit)
```

B.6 Performance Score

Code listing B.14: Calculating Performance Score

```
1 def calculate_score(mrei, resource_request_difference, ideal_target=95, wiggle_room
2     =4, mrei_weight=0.7, diff_weight=0.3):
3     mrei_difference = abs(mrei - ideal_target)
4     score_factor = abs(1 - (mrei_difference / ideal_target))
5
6     if resource_request_difference > 0:
7         resource_request_difference = 0
8
9     if mrei == ideal_target:
10        performance_score = 100
11    elif mrei_difference <= wiggle_room:
12        mrei_component = score_factor * mrei_weight * 100
13        diff_component = abs(resource_request_difference) * diff_weight
14        performance_score = mrei_component + diff_component
15    else:
16        mrei_component = score_factor **2 * mrei_weight * 100
17        print(mrei_component)
18        diff_component = abs(resource_request_difference) * diff_weight
19
20        performance_score = mrei_component + diff_component
21
22    # Clip the performance score to be between 0 and 100
23    performance_score = max(0, min(100, performance_score))
24
25    return performance_score
26
27 # Example usage
28 mrei = 100
29 resource_request_difference = +80
30 score = calculate_score(mrei, resource_request_difference)
```



```
31 print(f"Performance Score: {score}")
```