

BIG DATA CLOUD MIGRATION ASSESSMENT

Prepared for **Customer Name** at **Company Name**

November 29, 2022 - December 13, 2022

SUMMARY OF YOUR NEW CLOUD ENVIRONMENT

The estimated costs below assume on-demand pricing. Your actual costs might be significantly lower if you are able to take advantage of AWS pricing discounts.

\$11,409 estimated daily cost

\$342,282 estimated monthly cost

\$4,164,436 estimated annual cost

24% average resource utilization


Thank you for allowing Pepperdata to help you reduce your big data costs and achieve the benefits of migration to AWS. You can expect that a migration to AWS EMR would cost approximately **\$342,282 per month** based on your current workloads and on-demand instance pricing. For more information, please contact Name at name@companyname.com.

CLUSTER DETAILS FOR CLUSTER **Prod_Cluster**

November 29, 2022 - December 13, 2022

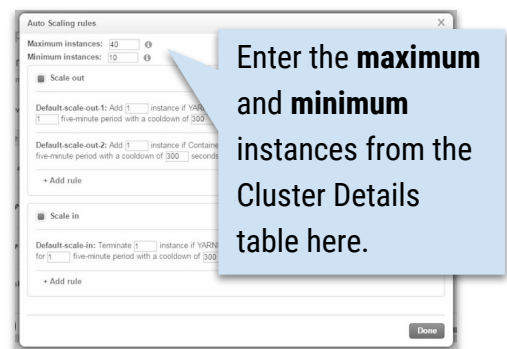
Queue Name	Recommended Instance Type	Minimum Instances	Maximum Instances	Projected Instance Hours (daily average)	Estimated Daily Cost	Estimated Annual Cost	Wasted Instance Hours (daily average)	Current Memory Utilization	Current CPU Utilization
root.users	r5a.16xlarge (cost: \$3.62 Memory: 512 GB 64 cores)	18	617	1,697.60	\$6,139	\$2,240,559	580.23	38%	37%
root.scale1	r5a.16xlarge (cost: \$3.62 Memory: 512 GB 64 cores)	1	318	722.81	\$2,614	\$953,998	424.73	43%	8%
root.queue3	r5a.16xlarge (cost: \$3.62 Memory: 512 GB 64 cores)	1	318	608.18	\$2,199	\$802,702	133.01	28%	24%
root.user8	m5a.24xlarge (cost: \$4.13 Memory: 394 GB 96 cores)	1	81	110.95	\$458	\$167,178	65.89	13%	5%
TOTAL / AVERAGE		21	1,334	3,139.55	\$11,409	\$4,164,436	1,203.86	31%	18%

ESTIMATED ANNUAL COST BY QUEUE (TOP 4 QUEUES)



OPTIMIZED AMAZON EMR AUTOSCALING PARAMETERS

Optimally chosen for your unique workloads



BIG DATA CLOUD MIGRATION ASSESSMENT

Prepared for **Customer Name** at **Company Name**

November 29, 2022 - December 13, 2022

TOP 10 USERS BY NUMBER OF APPS

User	MapReduce	Spark	Total
us9xfonc	6	6,915	6,921
ac9dusr	4,036	52	4,088
eu8dgg	1,412	61	1,473
jp9asb	1,160	43	1,203
acegusr		1,049	1,049
jws9usr	200	829	1,029
usxpusr	14	940	954
us9imta	544	98	642
co8ausr	167	405	572
im8tusr	38	485	523
All Others (155 users)	2,298	4,415	6,713
TOTAL	9,875	15,292	25,167

TOP 10 QUEUES BY NUMBER OF APPS

Queue	MapReduce	Spark	Total	Top User
root.user12	134	8,738	8,872	ur1089184
root.queue3	2,275	45	2,320	us8acm
root.br6	1,215	413	1,628	xg9usr
root.queue2	386	1,200	1,586	au8user
root.users	1,485	98	1,583	ca8eusr
root.scale1	1,416	70	1,486	xg8usr
root.pa18	41	1,269	1,310	au8uusr
root.user8	763	5	768	ug89dgdm
root.sc1		706	706	de8uice
root.cca	71	370	441	gb9ucpa
All Others (50 queues)	2,089	2,378	4,467	ab8ddusr
TOTAL	9,875	15,292	25,167	

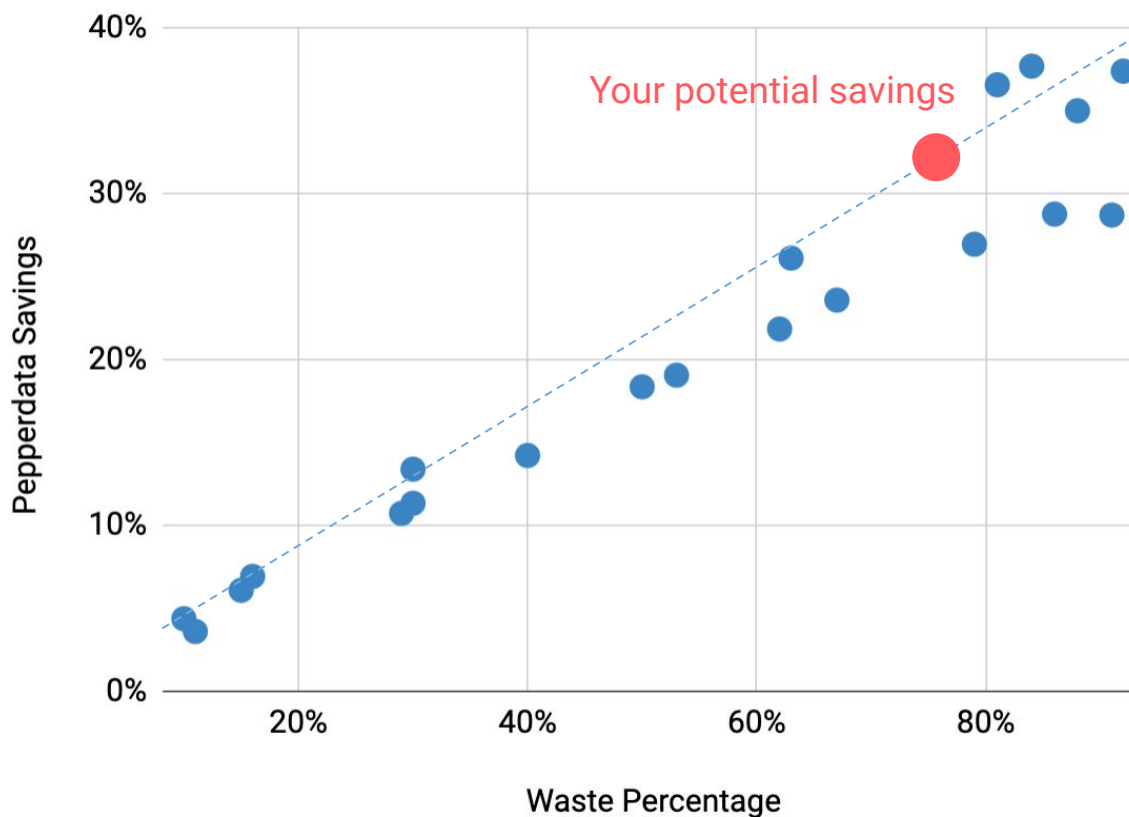
BIG DATA CLOUD MIGRATION ASSESSMENT

Prepared for **Customer Name** at **Company Name**

November 29, 2022 - December 13, 2022

PEPPERDATA OPTIMIZATION OPPORTUNITY

See how your environment compares with others optimized by Pepperdata



Once you've migrated your workloads to your new cloud environment, Pepperdata Capacity Optimizer allows you to reduce costs further. Capacity Optimizer analyzes the resource usage of each of your cluster's nodes in real time and uses machine learning to make thousands of resource allocation decisions per second. It identifies where more work can be done and adds tasks to nodes with available resources. The result: CPU, memory, and I/O resources are autonomously optimized to increase utilization, and waste is eliminated. Even the most experienced operator can't make manual configuration changes with Capacity Optimizer's level of precision and speed.


BIG DATA CLOUD MIGRATION ASSESSMENT

Prepared for **Customer Name** at **Company Name**


November 29, 2022 - December 13, 2022

PROJECTED MEMORY AND CPU ALLOCATION FOR **root.users**

Allocated Memory (GB) Over 14 Days



Allocated vCPU (Number of Cores) Over 14 Days



ALTERNATIVE INSTANCE TYPE RECOMMENDATIONS FOR **root.users**

Instance Type	Minimum Instances	Maximum Instances	Projected Instance Hours	Estimated Daily Cost	Estimated Monthly Cost	Estimated Annual Cost	Current Memory Utilization	Current CPU Utilization
TOP RECOMMENDATION								
r5a.16xlarge (cost: \$3.62 Memory: 512 GB 64 cores)	18	617	23,766	\$6,139	\$184,156	\$2,240,559	38%	37%
r6a.16xlarge (cost: \$3.63 Memory: 512 GB 64 cores)	18	617	23,766	\$6,160	\$184,807	\$2,248,490	38%	37%
x2idn.24xlarge (cost: \$10.00 Memory: 1536 GB 96 cores)	8	306	8,890	\$6,352	\$190,563	\$2,318,518	38%	37%
r5a.12xlarge (cost: \$2.71 Memory: 384 GB 48 cores)	25	843	32,807	\$6,355	\$190,656	\$2,319,648	38%	37%
x2idn.16xlarge (cost: \$6.67 Memory: 1024 GB 64 cores)	12	461	13,372	\$6,370	\$191,100	\$2,325,049	38%	37%


BIG DATA CLOUD MIGRATION ASSESSMENT

Prepared for **Customer Name** at **Company Name**


November 29, 2022 - December 13, 2022

PROJECTED MEMORY AND CPU ALLOCATION FOR **root.scale1**

Allocated Memory (GB) Over 14 Days



Allocated vCPU (Number of Cores) Over 14 Days



ALTERNATIVE INSTANCE TYPE RECOMMENDATIONS FOR **root.scale1**

Instance Type	Minimum Instances	Maximum Instances	Projected Instance Hours	Estimated Daily Cost	Estimated Monthly Cost	Estimated Annual Cost	Current Memory Utilization	Current CPU Utilization
TOP RECOMMENDATION								
r5a.16xlarge (cost: \$3.62 Memory: 512 GB 64 cores)	1	318	10,119	\$2,614	\$78,411	\$953,998	43%	8%
r6a.16xlarge (cost: \$3.63 Memory: 512 GB 64 cores)	1	318	10,119	\$2,623	\$78,688	\$957,375	43%	8%
r5a.24xlarge (cost: \$5.42 Memory: 768 GB 96 cores)	1	209	6,816	\$2,641	\$79,225	\$963,907	43%	8%
r6a.24xlarge (cost: \$5.44 Memory: 768 GB 96 cores)	1	209	6,816	\$2,650	\$79,506	\$967,319	43%	8%
r5a.12xlarge (cost: \$2.71 Memory: 384 GB 48 cores)	1	417	13,711	\$2,656	\$79,683	\$969,472	43%	8%


BIG DATA CLOUD MIGRATION ASSESSMENT

Prepared for **Customer Name** at **Company Name**


November 29, 2022 - December 13, 2022

PROJECTED MEMORY AND CPU ALLOCATION FOR **root.queue3**

Allocated Memory (GB) Over 14 Days



Allocated vCPU (Number of Cores) Over 14 Days



ALTERNATIVE INSTANCE TYPE RECOMMENDATIONS FOR **root.queue3**



Instance Type	Minimum Instances	Maximum Instances	Projected Instance Hours	Estimated Daily Cost	Estimated Monthly Cost	Estimated Annual Cost	Current Memory Utilization	Current CPU Utilization
TOP RECOMMENDATION								
r5a.16xlarge (cost: \$3.62 Memory: 512 GB 64 cores)	1	318	8,515	\$2,199	\$65,975	\$802,702	28%	24%
r6a.16xlarge (cost: \$3.63 Memory: 512 GB 64 cores)	1	318	8,515	\$2,207	\$66,209	\$805,543	28%	24%
r5a.24xlarge (cost: \$5.42 Memory: 768 GB 96 cores)	1	216	5,711	\$2,213	\$66,377	\$807,582	28%	24%
r6a.24xlarge (cost: \$5.44 Memory: 768 GB 96 cores)	1	216	5,711	\$2,220	\$66,612	\$810,441	28%	24%
r6a.32xlarge (cost: \$7.26 Memory: 1024 GB 128 cores)	1	159	4,344	\$2,252	\$67,564	\$822,027	28%	24%

BIG DATA CLOUD MIGRATION ASSESSMENT

Prepared for **Customer Name** at **Company Name**

November 29, 2022 - December 13, 2022

PROJECTED MEMORY AND CPU ALLOCATION FOR **root.user8**



ALTERNATIVE INSTANCE TYPE RECOMMENDATIONS FOR **root.user8**

Instance Type	Minimum Instances	Maximum Instances	Projected Instance Hours	Estimated Daily Cost	Estimated Monthly Cost	Estimated Annual Cost	Current Memory Utilization	Current CPU Utilization
TOP RECOMMENDATION m5a.24xlarge (cost: \$4.13 Memory: 384 GB 96 cores)	1	81	1,553	\$458	\$13,741	\$167,178	13%	5%
m6a.24xlarge (cost: \$4.15 Memory: 384 GB 96 cores)	1	81	1,553	\$460	\$13,805	\$167,955	13%	5%
m6a.32xlarge (cost: \$5.53 Memory: 512 GB 128 cores)	1	61	1,244	\$491	\$14,737	\$179,297	13%	5%
m5.24xlarge (cost: \$4.61 Memory: 384 GB 96 cores)	1	81	1,553	\$511	\$15,338	\$186,617	13%	5%
m5.metal (cost: \$4.61 Memory: 384 GB 96 cores)	1	81	1,553	\$511	\$15,338	\$186,617	13%	5%