

# Manage costs on Amazon EMR

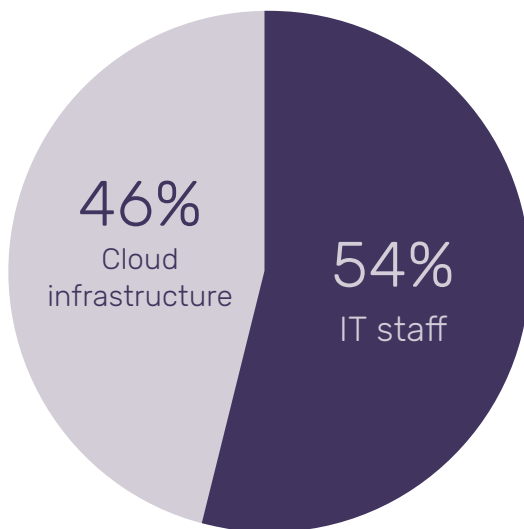
Operationalizing big data to optimize and automate modern applications and infrastructure.



# Main cost drivers

Amazon EMR is a service that makes it easy to process large amounts of data efficiently. It uses Hadoop processing combined with several AWS products to do such tasks as data mining, log file analysis, machine learning and scientific simulation through a managed service. It has become very popular as a cloud alternative to physical hosting of complex infrastructure. However, customers are finding unexpected costs eating into their cloud budget. Furthermore, lack of visibility to root cause and general inefficiency is costing organizations thousands, if not millions in operating their Amazon EMR environment.

The following chart shows a breakdown of typical Total Cost of Ownership for Amazon EMR project costs:



Controlling IT Staff productivity costs (quick troubleshooting, meeting SLAs) are just as important as controlling cloud infrastructure costs (compute, storage, networking) in Amazon EMR. To address both human capital and infrastructure costs, the following issues are the largest contributors to cost overruns in Amazon EMR:

1. **Poorly performing or failed jobs** - EMR environments often have hundreds, if not thousands of jobs running simultaneously. Identifying poorly performing jobs often occur after resources are consumed.
2. **Getting to root cause analysis** - EMR administrators sift through thousands of lines of logs to identify and remediate the root cause of failed or underperforming jobs.
3. **Missed service level agreements** - Missed SLAs for business applications cost an organization real dollars in missed opportunity or financial penalties.
4. **DevOps/IT Ops blame game** - Failure to quickly identify root cause will lead to reactionary finger-pointing between the app developers (“it’s the cluster”) and IT Ops (“it’s the code”).
5. **Chargeback/showback** - Inadequate tracking of which department or user may be abusing cluster resources perpetuate cost overruns.

The issues in DataOps have been present for some time, but we are now better assessing the costs of these inefficiencies in big data projects. These issues contribute to direct organizational costs attributed to developing and operating an Amazon EMR environment. As a result, the true costs of missed SLAs could have vast financial impacts on your business:

Consequences of missed SLAs	
Banking	Fraudulent transactions in banking not discovered in time, leading to millions in theft and fines.
Healthcare	Patient profiles not accurately described, leading to lapses in care and millions in additional healthcare costs.
Retail	Customer demand not analyzed for a particular product, overestimating inventory leading to millions in waste.
Manufacturing	Equipment failure not detected with accuracy, leading to costly maintenance calls.

# Tuning Amazon EMR requires awareness of resources

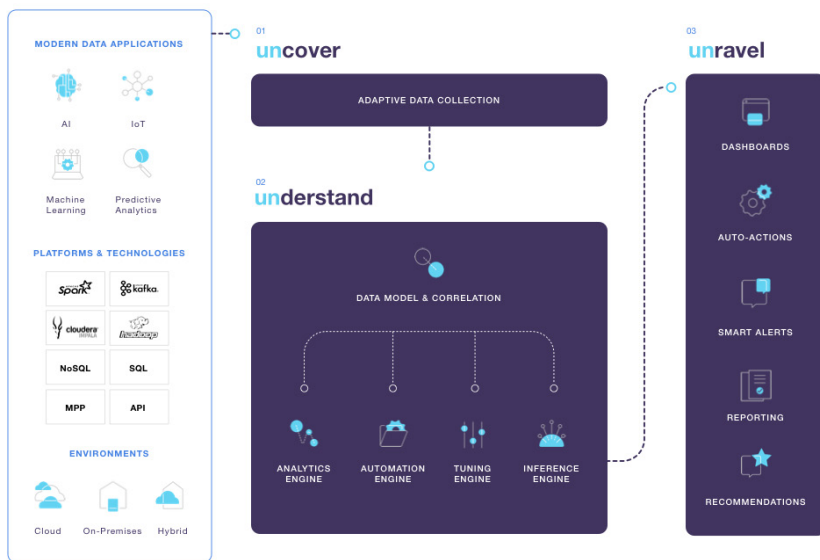
There are a few money saving techniques which will help avoid cost overruns:

- **Choose Proper AWS EC2 topology** - Properly assigning the correct AWS EC2 instance types and number of nodes is essential to control spending.
- **Choose proper storage and networking** - Decide on the networking configuration (is public IP necessary?) as well as anticipating optimal storage (S3 types).
- **Code optimization** - Inefficient Spark or Hive code in your Amazon EMR application can kill your performance or cause failures - tight collaboration between DevOps and ITops is essential to provision right level of resources.
- **Workload identification** - Not all Amazon EMR jobs are created equal. Data warehousing will have more persistent resource requirements (and cost) than ephemeral data engineering jobs, for example. Anticipate these needs and adjust accordingly.
- **Manual instance matching** - developers are able to choose from a variety of workloads (data analytics, data engineering, data engineering light) as well as Premium/Standard levels within each workload.
- **Unused compute elimination** - Unchecked rogue resources can be a large contributor to waste. Users must understand when auto termination is warranted.

These techniques require manual data collection, but can be effective for managers versed in log collection techniques for Spark, Hive, MapReduce and Kafka.

# Unravel's solution

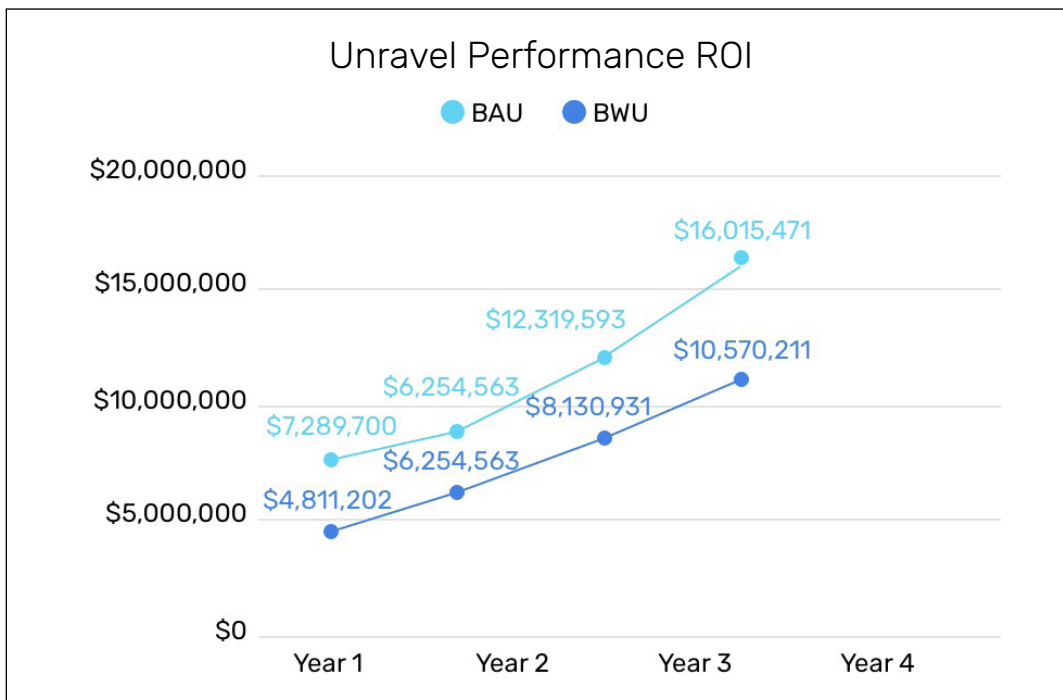
Alternatively, Unravel provides a turn-key solution for cost assurance.



With Unravel	Without Unravel
<b>Time to resolution:</b> Hours to minutes	<b>Time to resolution:</b> Days to weeks
<b>Required Infrastructure:</b> S3 - \$100K/year EC2 - \$200K/year EMR - \$100K/year	<b>Required Infrastructure:</b> S3 - \$150K/year EC2 - \$500K/year EMR - \$200K/year
<b>Required Team:</b> Big Data Engineers - 10 hours/week	<b>Required Team:</b> Big Data Engineers - 40 hrs/week
<b>DevOps:</b> 5 hours/week	<b>DevOps:</b> 10 hours/week

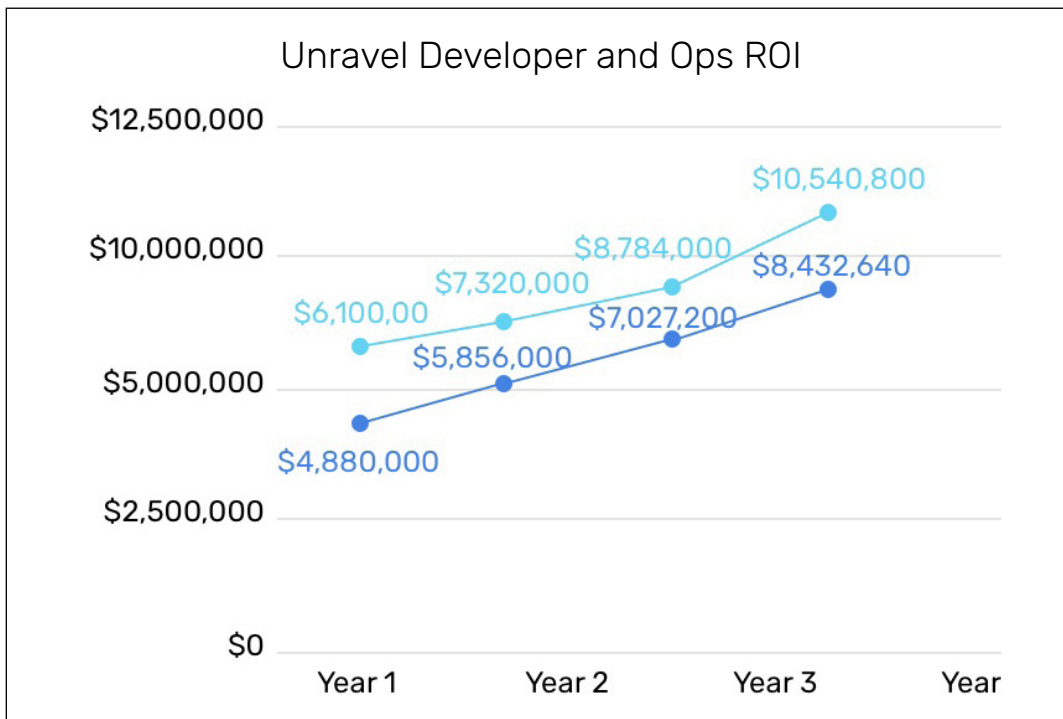
# Example: Unravel banking customer TCO

The following diagram illustrates a real life banking customer's infrastructure savings using Unravel:



**Unravel will drive \$9MM in infrastructure savings over 3 years.**  
Assumes 20% annual growth rate.

The following diagram illustrates a real life banking customer’s labor savings using Unravel:



***Unravel will drive \$6.5MM in human capital savings over 4 years.***  
*Assumes 20% employee growth.*

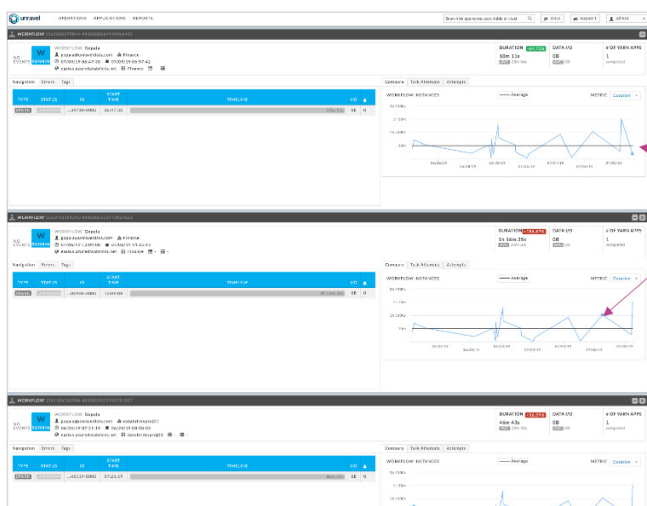
Given the growth rate of data, and the appetite for enterprise organizations to analyze this data in real time, companies are facing increasing costs in infrastructure (both on-premises and cloud) and human capital to maintain analytics projects.

Unravel presents an elegant solution to provide cost assurance in Amazon EMR environments, automatically collecting and correlating jobs data in context, helping companies pinpoint issues, often before they happen.

# Unravel techniques for Amazon EMR tuning

## 1. Poorly performing or failed jobs

**Visualize jobs and job runs:** track individual jobs to assess performance improvements.



- Easily compare runs of a job and their KPIs
- Find out the reasons for difference in perf

Various runs of a job

Easy to compare various attributes – Duration, I/O across runs of the same job etc.

Run ID	Start Time	Completed	Duration
Run 187	2019-07-08 14:48:17 PST	Manually	5m 52s
Run 188	2019-07-08 12:43:37 PST	Manually	8m 47s
Run 189	2019-07-08 13:01:23 PST	Manually	8m 17s
Run 194	2019-07-08 13:08:23 PST	Manually	8m 14s
Run 195	2019-07-08 08:40:20 PST	Manually	4m 57s
Run 192	2019-07-08 09:04:07 PST	Manually	4m 57s
Run 181	2019-07-08 05:41:15 PST	Manually	4m 46s
Run 180	2019-07-07 07:17:02 PST	Manually	4m 37s
Run 179	2019-07-07 06:08:20 PST	Manually	7m 48s
Run 178	2019-07-07 04:30:23 PST	Manually	4m 28s
Run 177	2019-07-07 03:08:20 PST	Manually	4m 45s
Run 176	2019-07-07 02:48:19 PST	Manually	10m 17s
Run 175	2019-07-07 02:48:19 PST	Manually	4m 37s
Run 174	2019-07-07 01:04:07 PST	Manually	4m 45s
Run 173	2019-07-07 01:01:18 PST	Manually	4m 47s
Run 172	2019-07-07 01:18:30 PST	Manually	4m 47s
Run 171	2019-07-07 00:00:00 PST	Manually	4m 38s
Run 170	2019-07-07 00:24:31 PST	Manually	4m 37s
Run 169	2019-07-07 00:24:31 PST	Manually	4m 38s
Run 168	2019-07-08 14:48:17 PST	Manually	4m 38s
Run 167	2019-07-08 14:48:17 PST	Manually	4m 38s

## 2. Getting to root cause analysis

**Root cause analysis:** if we see a resource having issues we can use point of time KPIs to identify the state of the applications at the point of failure.

**EVENTS PANEL**

Efficiency 3 Application Failure 1

SPARK app-20190703071859-0000

**EXECUTOR FAILED WITH OUTFORMEMORY ERROR**

2 executors in the application have failed with Java OutOfMemory error

Check the [usedHeap, maxHeap, committedHeap, totalPhysicalMemory, availableMemory, vmRss, gcLoad] metrics for executors [executor-2, executor-0] on the Resource Usage Lab

If application is running in non-local mode, then increase the JVM heap size via the --executor-memory parameter or the spark.executor.memory property

Specific reason for failure extracted and presented - RCA is fast and easy. No need to dig through logs

Attempt 1

Navigation	Execution	Gantt Chart	Errors	Logs	Conf	Tags
WARNING	07/03/19 00:19:05		ThrottledLogger\$:	Failed to load user identity when		
WARNING	07/03/19 00:19:05		EC2MetadataUtils:	Unable to retrieve the requested		
FATAL / EXCEPTION	07/03/19 00:19:40		java.lang.OutOfMemoryError	java.lang.OutOfMemoryError: Gi		
ERROR	07/03/19 00:19:41		TaskSchedulerImpl:	Lost executor 0 on 10.139.64.6: I WARN messages.		
ERROR	07/03/19 00:20:54		TaskSetManager:	Task 4 in stage 0.0 failed 4 times		
FATAL / EXCEPTION	07/03/19 00:20:53		GC overhead limit exceeded	Lost task 4.2 in stage 0.0 (TID 6, <mini>AbstractStringBuilder.jav: (StringBuilder.scala:47) at scala.com.unraveldata.spark.test.Spari at com.unraveldata.spark.test.Spari at scala.collection.TraversableLike scala.collection.TraversableLike scala.com.unraveldata.spark.test.Spari com.unraveldata.spark.test.Spari com.unraveldata.spark.test.Spari		

List of errors and relevant information extracted and surfaced

SPARK spark

Duration: 1m 56s | Data I/O: 0 B | # of Jobs: 1 | # of Stages: 1

Attempt 1

Navigation: Execution | Gantt Chart | Errors | Logs | Conf | Tags

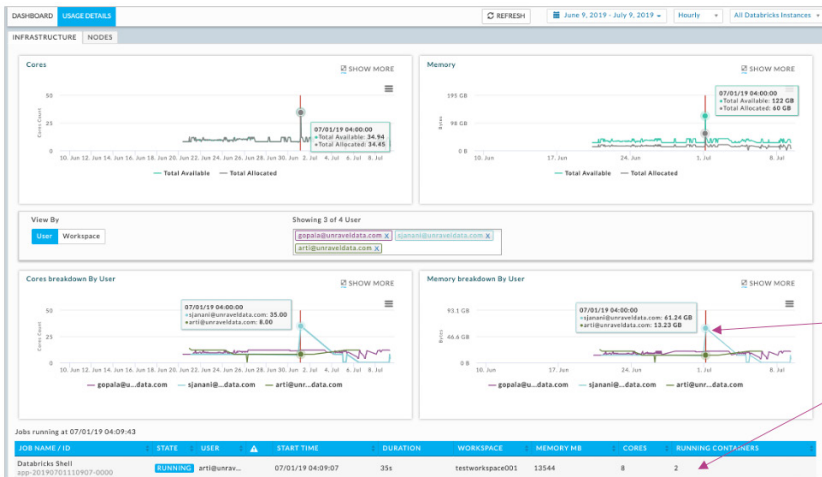
Program: Task Attempts | Resource

Copy

Relevant logs retrieved and made easily available

### 3. Missed service level agreements

**Single pane of glass view:** ensure maximum SLAs by pinpointing possible failures before they happen.

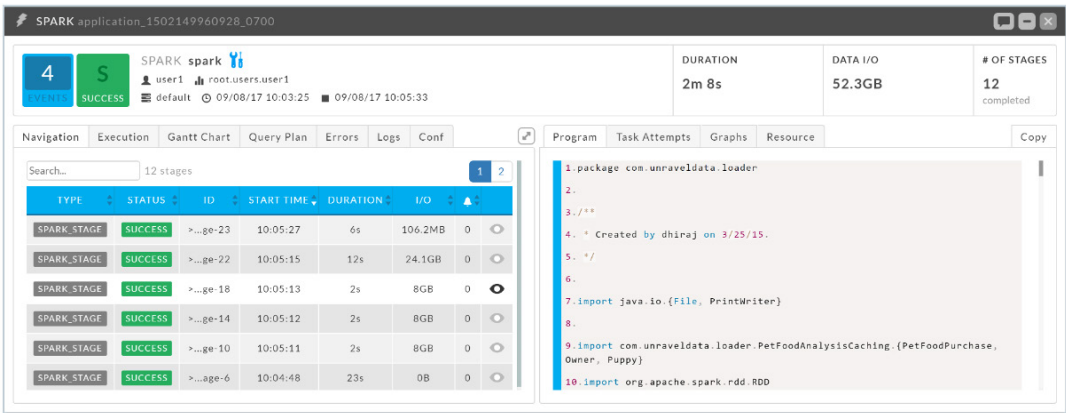


View Resource Usage, Jobs, Users across Databricks Instances & Workspaces

Sudden rise seen in resource usage  
Identifying the list of job runs that caused that peak

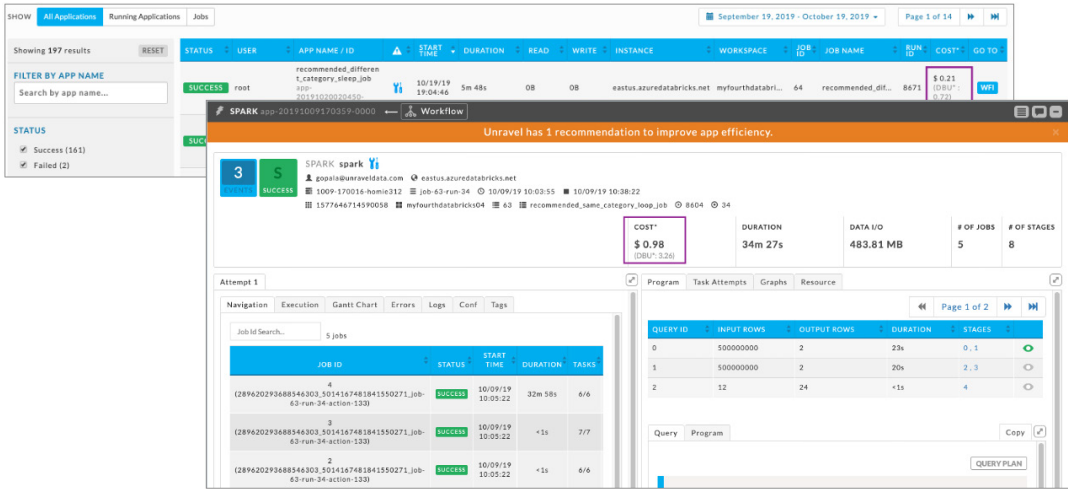
### 4. DevOps/ITOps blame game

**AI-Driven recommendations:** pinpoint whether the issue resides in the code or the cluster configuration via automatic recommendations.



### 5. Chargeback/showback

**Per application costs:** Identify resource wasters by application or by user to drive accountability for responsible EMR resource use.



Interested in learning more? Contact us at [hello@unraveldata.com](mailto:hello@unraveldata.com)

#### About Unravel

Unravel radically simplifies the way businesses understand and optimize the performance of their modern data applications – and the complex pipelines that power those applications. Providing a unified view across the entire stack, Unravel’s AI-powered data operations platform leverages AI, machine learning, and advanced analytics to offer actionable recommendations and automation for tuning, troubleshooting, and improving performance – both today and tomorrow.

By operationalizing how you do data, Unravel’s solutions support modern big data leaders, including Kaiser Permanente, TIAA, Adobe, Deutsche Bank, Wayfair, and Neustar. The company is headquartered in Palo Alto, California, and is backed by Menlo Ventures, GGV Capital, M12, Data Elite Ventures, and Jyoti Bansal. To learn more, visit [unraveldata.com](http://unraveldata.com).

