

Comprehensive Chargeback Reporting: Calculate Accurate Costs for Multi-Tenant Hadoop

Pepperdata's new chargeback reporting feature allows you to view the costs of using shared cluster resources — like memory or CPU — over a specified period of time. If you are building a multi-tenant or multi-workload Hadoop deployment, you can now easily view consumption by user, job, or department to quickly learn which are the most costly and most resource-intensive, calculate ROI, and make strategic decisions around IT investment. If you are deploying internal Hadoop-as-a-service, chargeback reporting will allow you to build the business processes you need to allocate costs across departments or workgroups.

user	physical memory used		cpu used	
	byte-hours ↓ ₁ ⁹	cost	cpu-hours	cost
etl	42358.59G	\$4231.12	45257.72	\$4971.97
forecast	35471.18G	\$3519.66	49331.40	\$5492.65
hbase	25009.24G	\$2473.46	6183.64	\$734.48
analytics	21597.18G	\$2159.73	34949.58	\$3229.03
dev	9188.56G	\$918.85	6046.20	\$657.93
reporting	5084.35G	\$549.65	11760.50	\$1248.19

Key Benefits

- View physical memory and CPU usage by user, job, or workflow
- Quickly see which users, jobs, or workflows use the most resources — and what this costs
- View reports from any Internet-accessible device
- Get real-time visibility into resource consumption
- Make more informed strategic decisions around IT purchasing and planning
- Use without needing any additional installation
- Configure exact per GB-hour memory and per CPU core hour costs for accurate cost analysis
- Enhance cost awareness across business units
- Export cost data into CSV with one mouse click

About Pepperdata

Pepperdata is the leader in real-time cluster optimization. Our software brings predictability to multi-tenant, production Hadoop environments for the first time. While Hadoop and YARN focus on getting jobs running on the cluster, Pepperdata monitors and controls hardware usage in real time so that customers can meet SLAs, increase throughput, and improve visibility.

Pepperdata software is trusted on over 10,000 production nodes in firms of all sizes — including 3 Fortune 50 companies.

