NEW
RESOURCES
CONSULTING

USE CASE

Leveraging Dynamic Scaling in Oracle Exadata

Challenge

A public sector agency was managing significant system processing loads that occurred during peak windows—such as timesheet entry periods on Monday mornings and Friday afternoons, overnight batch processing, and overnight data warehouse ETLs. To maintain acceptable performance during these spikes, infrastructure was sized for peak demand, resulting in costly overprovisioning.

Outside of these peak periods, systems were operating at less than 30% CPU utilization, creating inefficiencies and unnecessary expense. Facing a mandate to reduce overall spend by 10 – 20%, the agency engaged New Resources Consulting (NRC)'s Enterprise Solutions Group to optimize infrastructure costs without impacting performance.

Approach

NRC partnered with the client's DBA team and Application Manager to conduct a detailed performance analysis, identify optimization opportunities, and design a dynamic scaling strategy that leveraged existing Oracle Exadata capabilities. Through performance analysis and ongoing advisory support, the team aligned on peak processing requirements and defined acceptable minimum capacity levels for day-to-day operations.

The engagement focused on leveraging existing platform capabilities, validating changes through monitoring in non-production and production environments, and ensuring no degradation in system performance, with an overarching goal of reducing operating costs by a minimum of 10%.

1000 N Water St, Ste 950
Milwaukee, WI 53202

414.289.7960

info@newresources.com

newresources.com

Solution

Working collaboratively with the client's DBA team, NRC designed and guided the implementation of dynamic infrastructure scaling using Oracle Exadata capabilities.

The approach included:

- Automating resource scaling based on workload demand using modified Oracle-provided scripts
- Introducing monitoring tools to track utilization and measure cost savings
- Continuously analyzing system performance to ensure stability during scaling adjustments

NRC provided ongoing advisory support throughout the engagement, defining peak processing requirements, establishing acceptable minimum capacity thresholds for day-to-day operations, and validating changes through rigorous testing in non-production and production environments.

This approach allowed the system to scale up during peak demand and scale down during off-peak periods—without impacting end users or application performance.

Results

The solution delivered measurable cost savings and operational improvements:

- **Cost Efficiency:** By dynamically aligning infrastructure capacity with actual demand, the organization achieved a 15% reduction in infrastructure costs—exceeding initial savings targets—while eliminating unnecessary spend during non-peak periods.
- **Performance Stability:** Despite reduced resource utilization, system performance remained consistent, ensuring no negative impact on end users or application responsiveness.
- **Operational Visibility:** Enhanced monitoring and utilization insights provided greater transparency into system performance, enabling more proactive and ongoing optimization efforts.
- **Code Optimization:** With infrastructure costs now tied more directly to usage, inefficient custom code became more visible, driving improved performance tuning and more efficient development practices.
- **Scalable Optimization Model:** This approach demonstrates that cost savings are not limited to a single platform. By leveraging flexible, usage-based infrastructure, the model can be extended across other virtualization and cloud environments to achieve similar efficiencies.

Overall, the organization achieved a more efficient, cost-effective infrastructure model while maintaining performance standards. The success of this engagement has positioned NRC as a trusted advisor for ongoing optimization initiatives across the agency's enterprise infrastructure.