



Automated Optimization PowerVM on AIX

Most organizations combat risk in virtualized infrastructure by significantly over-provisioning hardware. Excess capacity is a cost that is seen as a necessity because of the complexity of these new school environments. But this inefficiency can be avoided by optimizing workload placements and right-sizing LPAR allocations to simultaneously combat both risk and capacity waste.

Densify is a cloud optimization technology focused on tackling the skyrocketing costs of cloud adoption. Delivered as a service, its patented technology uses machine learning to automatically learn your cloud usage patterns 24/7, and proactively recommends what changes you should make and results to expect. Densify's cloud migration analysis enables you to plan the fastest, safest and most cost-effective path to the cloud.

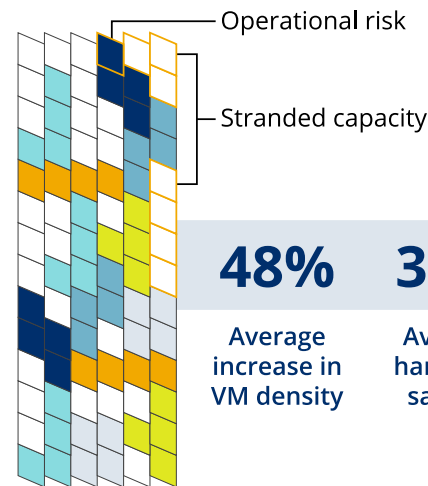
Optimize LPAR Placements to Increase Efficiency

Densify's predictive analytics optimize LPAR placements considering all technical and business requirements and detailed utilization patterns and personalities to significantly increase density while reducing the risk of resource contention. Leveraging patented analytics, Densify uniquely dove-tails workloads, optimally fitting LPARs together on the hosting frame. Think of it like a game of Tetris®, poorly placed workloads strand capacity and can lead to resource contention, while optimally placed workloads make the best possible use of the physical infrastructure. The analytics powering Densify increases LPAR density and reduces hardware requirements.

Unique Insight into Your Infrastructure

Densify uniquely dovetail workloads like in the game of Tetris® to optimally increase density. This means your applications run better and need less infrastructure, both on-prem and in the cloud.

Before Densify



48%

Average increase in VM density


33%

Average hardware savings

55%

Average software license savings

After Densify



As environments grow and as workloads change, what was once optimal may no longer be. Densify.com predicts when changes in the environment are required and rebalances workloads as conditions change. Unlike a load balancer, Densify.com's placements take predictive view, analyzing trends over time to avoid performance issues and dramatically reduce LPAR motioning and volatility.

Daily actions enable you to increase efficiency by reclaiming resources and optimizing use of existing capacity such as:

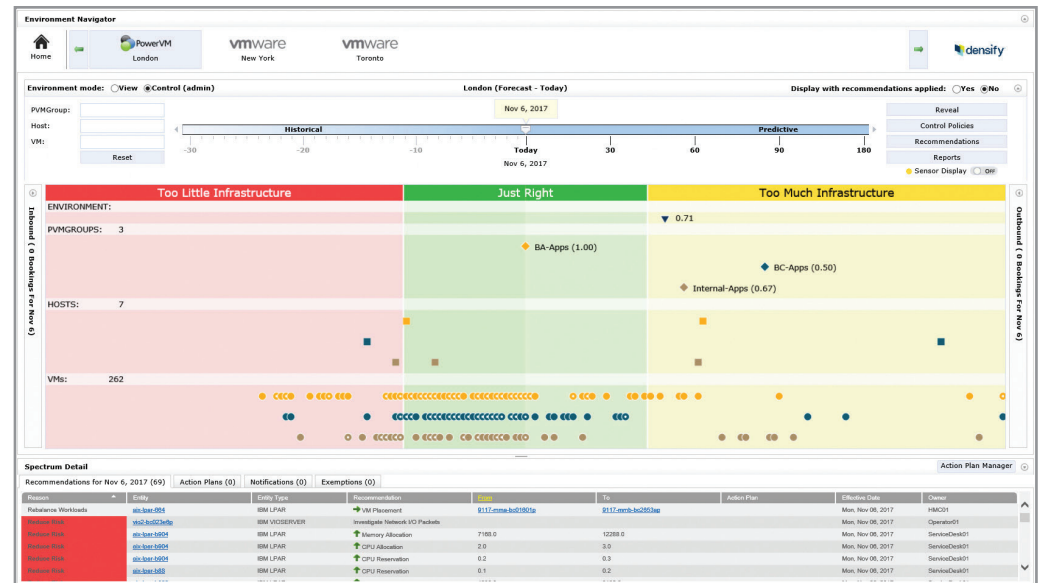
- **Over-allocated LPARs** – Over-sized LPARs are a common area of waste. Densify's allocation recommendations ensure that LPARs are sized correctly and consider minimum/ maximum allocation and entitlement values for LPAR CPU configuration and also makes right-sizing memory recommendations.
- **Inefficient workload placements** are resolved with precise placements that safely increase density and fit the workloads onto the minimum amount of infrastructure.

Predictive Analytics to De-Risk

Densify's predictive analytics optimize resource allocations and placements to proactively avoid operational risk. Its advanced algorithms predict resource contention and workload patterns and provide prescriptive recommendations that right-size and optimally place workloads and avoid issues.

Detailed reports and prescriptive, automatable recommendations address:

- **Under-allocated LPARs** enabling you to right-size LPARs whether you are using a standard instance catalog for clouds or custom LPAR allocations.
- **Hosting frame imbalances** with rebalancing to ensure each workload gets the resources it needs and changes are made in advance to avoid performance issues.
- **Inadequate capacity at the server farm and environment levels** through visibility into resourcing status at a farm or environment



Densify also optimizes:

vmware® PowerVM IBM aws Google CloudPlatform Microsoft Azure

