



Case Study

Deploying HPC Service at Leading Aerospace Company



CHALLENGES

Fortune

50

Aerospace Client

Consolidate

16

HPC sites

Reliable service for

1000s

of engineers

CBT architects co-developed, centralized High-Performance Computing (HPC) service, providing pay-as-you-go HPC resources for thousands of engineers.

Background

CBT is a premier, woman-owned technology solutions provider with **extensive experience designing and deploying enterprise-class solutions** for clients across the globe. By combining best-of-breed technologies, CBT provides a balance of **services and solutions in the areas of HPC & Analytics, Industrial IoT, Hybrid IT and IT Supply Chain Optimization** to augment our customers with the necessary tools to excel in the global marketplace.

A perfect example of how CBT delivers value to its customers is the following **HPC service success story** at one of the world's largest aerospace companies. As the leading manufacturer of commercial jetliners, as well as defense, space and security systems, this company has multiple large engineering organizations with extremely high demands for HPC resources.



Case Study

Deploying HPC Service at Leading Aerospace Company

“

We pulled together these pockets of HPC and saved a ton of money.

CUSTOMER CIO

Problem

Over time, and over several mergers and acquisitions, this aerospace company's different sub-business units had ballooned to **more than 16 independent HPC sites**. Each being used by thousands of engineers spread across continental USA.

Although utilizing their equipment in similar capacities, each site's HPC team built their systems to suit their particular preferences and expertise. The **environmental differences** included multiple vendors and interconnect choices, asset age, and even encompassed varying management styles and accounting rules. The physical distance between sites and the lack of infrastructure availability added so much latency to each transaction that sharing resources, or migrating workloads between sites, was not an option.

Each HPC site ran a variety of **highly demanding, computer-aided engineering applications** such as Computational Fluid Dynamics (CFD), structures and materials, computational electromagnetics, molecular modeling, and multiple other simulation applications. The size and scope of any one job, and the ability to run multiple jobs concurrently, was restricted by the limited HPC resources available on these isolated small sites.

While the flexibility for each site to manage their own resources may seem useful on the surface, the reality was, it led to **massive inefficiencies** and **unnecessary expenses**. At any one time, you could see a mix of sites with too much work (over-subscription), or not enough work to keep the entire system running at full blast (underutilization or white space).

Additionally, HPC application software licenses are exceptionally expensive, meaning **inefficient usage comes at a high price**. The goal is to maximize the number of hours per month that each license is used – that is, run jobs more efficiently, without extended downtime, to get the most output from a fixed number of licenses. In this case, licenses were isolated to each site and usage was restricted by the same over-subscription and underutilization issues as the HPC infrastructure.

Moving HPC workloads to a public cloud like AWS was considered, but ultimately ruled out as a solution based on availability, cost, risk, and security regulatory compliance requirements.



Case Study

Deploying HPC Service at Leading Aerospace Company

Solution

To ensure the aerospace company's engineers got the compute resources they needed while improving cost efficiency, it was decided they should build a **centralized HPC service** across (2) sites, thus providing a larger shared HPC resource pool and standardizing the way HPC was provided and consumed for multiple business units.

First, a **High-Performance Computing Council (HPCC)** was established to gather resource requirements and supporting business cases, provide capacity planning information, and be the oversight in resolving conflicts for emergency (unforecasted) workloads. Based on the customer's HPCC's requirements and business case, CBT worked with the Operations Staff to begin architecting the final solution.

Next, CBT established a **series of on-site technology briefing sessions** for the customer and all the vendors involved. These customized briefings brought customer stakeholders together to quickly understand new solution capabilities and technology roadmaps, therefore allowing **collaboration on integrated design work** to begin much sooner. CBT also arranged for **on-site evaluation equipment** and **remote testing facilities**, including early adopter programs with pre-production equipment and software. These CBT resources allowed solution design based on real hands-on experiences, facts, and test data, rather than stale paper studies and catalog data.

Finally, CBT helped define and deliver a **complete bill of materials (BOM)** for a centralized HPC environment, including state-of-the-art storage and compute technologies, low latency and high bandwidth interconnects (InfiniBand) – all delivered via internal private cloud and Virtual Desktop Infrastructure (VDI) technologies for optimized ROI and TCO. CBT, as the overall solution integrator, had a leading role in **meeting the technical requirements, staying within the defined budgets, meeting delivery schedules, and delivering a fully functional service** before the targeted start date.



Case Study

Deploying HPC Service at Leading Aerospace Company

RESULTS

1

Centralized HPC Service

3-4x

more work per license

FUTURE

From 2010-2020

~14k%

increase in data from airplanes

Results

By consolidating 16 islands of HPC into one **centralized enterprise resource**, CBT **helped unify the customer's operations, centralize workload management with a shared scheduler, and implement a chargeback system** (based on cost per node-hours) to stabilize resource allocation and assist with company-wide forecast management.

With **combined workloads** and **capital investment** from all 16 sites, the aerospace customer could purchase a much more powerful and cost-effective system than any one of the smaller sites could have afforded. With a vastly faster system based on higher-end parts, efficiencies are increased by running jobs faster and getting them off the system sooner. This provides more usable hours for each software license, resulting in three to four times more work per licenses, at the same expense.

Utilizing their new HPC environment as a "not for profit" service, the customer charges business units based on consumption (with a break even cost recovery pool). Given that there are none of the additional fees or margins normally associated with external cloud services, this system ensures business units maximize the amount of HPC work per dollar spent.

Looking Forward

Advanced systems with increased power and efficiency, like this HPC service, open the door for **cross-departmental collaboration** on some exciting new projects. For example, by combining the expertise and resources of the HPC and Analytics teams, the aerospace customer hopes to employ advanced concepts such as finite model theory to improve the precision of their structural designs. They could run big data analytics and increase the resolution of CFD models for improved aerodynamics. This is especially important as the aerospace company integrates data collected by airplanes in flight and feeds it back into the design and development process.

The aerospace company's CIO and senior vice president of information and analytics, comments, *"Between 2010 and 2020 we estimate there will be a **14,000% increase** in the amount of data coming off modern airplanes. We see this data as a way to **improve the way we design, engineer, and manufacture by providing a feedback loop of all that data back into the company processes.**"*