



The Drive to Greater Data Center Efficiency



The Drive to Greater Data Center Efficiency

In today's business climate, most Fortune 1000 companies face pressure to reduce or maintain flat IT costs, while delivering new business applications and services at an ever increasing pace. Though data centers are bursting at the seams, there is little appetite to build out more facilities (it can cost around \$100M per data center for a large enterprise). Yet IT is at a steep disadvantage: how can they quickly identify opportunities for cost savings, without an accurate view of their applications and infrastructure and or a sense of how that infrastructure relates to the delivery of business services.

This lack of visibility into the existing infrastructure, and how it is being used to support the business, results in inefficient data center resource planning. Too often, more infrastructure than is necessary is deployed to meet business objectives. The result is massive inefficiency. In the average large data center:

- 10% of servers perform no current usable function
- 2.5%-5% of servers are marked as decommissioned in the fixed asset register – yet are still operational
- Most servers run at 10-15% utilisation

Each of these servers has associated administration, software licence, facilities, power and cooling costs that all contribute to the rising IT costs of just 'keeping the lights on'. The deteriorating economy, the pressure to be 'green', increasingly high expectations for business service delivery all converge to underscore the need to optimize data center infrastructure. The ultimate goal? Ensure IT can continue to support business-as-usual services – as well as provide room to innovate – without growing in size and cost-base.

The good news is that there is a significant and immediate opportunity for cost saving by taking out the inefficiencies that exist – starting with an understanding of how each server is used, which business applications it supports, and therefore what the value to the business actually is. In this way unused, unnecessary and inefficient hardware can be weeded out or refreshed, and software licenses reclaimed while ensuring compliance in the year-end "true-up" exercise.

Common Objectives of Data Center Optimization

Data Center Optimization projects can be triggered by one of several business objectives, including:

- Storage and server virtualisation projects
- Technology refresh to replace end-of-life, poorly supported or expensive to maintain infrastructure
- Provision of disaster recovery capability

- Cost reduction programs
- Data center space constraints (driven by available real-estate as well as thermal/size/power constraints)
- Mergers and acquisitions
- Organizational policy and strategy – including regulatory considerations
- Standardization (for licensing savings, skill set optimization and reducing complexity)
- Software license reclamation
- Disposal of real estate
- Removal of an outsourcer

Truly proactive IT organizations are realizing the need to get ahead of the curve and start to continuously optimize data center resources as a 'business as usual' activity.

The Benefits of Data Center Optimization

Data center optimization can deliver significant benefits if executed effectively:

- Lower costs through more efficient use of infrastructure with associated lower maintenance costs. This includes reductions in: physical hardware (including virtualization); software licensing; personnel requirements and real-estate.
- Better management of operational risk including concentration risk, resiliency and BCP.
- Increased flexibility to support business requirements as there is more head room to deploy new services – often, overstretched datacenters only allow new infrastructure in exchange for removing old infrastructure.
- Improved service levels including fewer change-related outages
- Tightened security through improved standardization and centralized management.
- Simplified DR by reducing the number of data center locations, and standardizing server builds, planning DR becomes a far simpler process with fewer moving parts to consider.
- Reduced administrative overhead by centralizing management, and reducing breadth of skill set requirements.

These benefits would be hard to achieve today given the current, limited visibility into the IT environment. Often, even knowing where to start can be daunting, and audit processes can take person-years to achieve. Lack of understanding of the dependencies between software, applications and infrastructure cause downstream issues, including changes in scope and resultant failure to meet deadlines and budgets.

Types of optimization

There are five main types of data center optimization – one or more may be performed at any one time. The more that are attempted simultaneously, the greater the risk.

- **Refresh** – replacement with a newer, more cost-efficient resource
- **Consolidate** - the migration from several data centers, servers or instances of software into fewer.
- **Retire** - decommission old or obsolete equipment, software and business applications.
- **Move** – forklift move of infrastructure to new facilities, or swinging from an old to an identical new environment.
- **Virtualize** – replacing physical servers with virtual servers
- **Reclaim** – reclaim unused software licences for the annual true-up exercise

Every Journey Starts with a Map

Many firms embark on ambitious data center optimization projects, with these benefits in mind, only to realize they are lacking fundamental data about their environment. In the earliest planning stages of the project, it is critical to gain a complete map of the data center resources and the application dependencies - fast. It is also important to realize that no data center stands still – changes will occur all the time, causing configuration drift from that baseline. Even when dependencies are fully documented at the start of the project, by the time the move for an application or system finally occurs, it's highly likely that some portion of the topology or dependencies has changed.

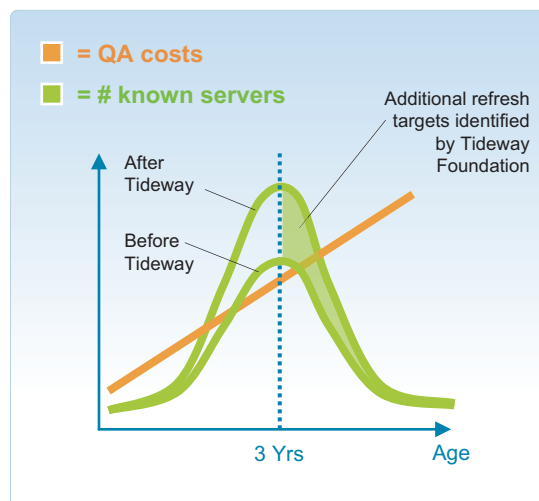
Today, physical resource information is often gathered manually, increasing the likelihood of error and the overall time required. Application dependencies tend to be gathered through interviews with the application owners, which distracts them from the day to day responsibilities. When built in this way, the overall data quality of the system of record from which planning decisions are made is only around 50-70%.

Not having an accurate map of application dependencies in the data center can cause painful outages when moves occur – often not only to the application being moved, but to related applications too. A failed move is many times as expensive to correct as the cost of the move itself – again causing additional costs and time overruns to the original project schedule. The risks of application downtime caused by lack of understanding of dependencies causes many organizations to insist on only performing 'like-for-like' application swings when performing a data center move. This minimizes the moving parts that could fail, but misses the opportunity to refresh and therefore optimize.

Stranded or “forgotten” assets litter datacenters. Across dozens of implementations, Tideway consistently finds that 2.5-5% of servers still active in the data center are marked as decommissioned in the fixed asset register! These are usually fairly old machines, which are therefore likely to be especially large and inefficient. Compound that with another alarming statistic: up to 10% of servers in the data center have no current usable function – though they do manage to consume large amounts of electricity, generate heat, and occupy valuable data center floor space. These represent a great opportunity to retire for immediate benefit – provided you can find them.

As data center infrastructure gets older, not only is it larger and less efficient than newer infrastructure, but the costs of quality assurance rise – failures are more likely, maintenance is more expensive and skillsets are potentially rarer (read – more expensive). Refreshing this infrastructure is an obvious way to improve the situation, however finding refresh targets can be problematic – the older a machine is, the less likely the system of record is to have an accurate representation.

What is clear is that the processes that give us visibility of the hardware, software and systems across our IT estates are ripe for an overhaul – they are manual at best, and due to the high rate of change, always produce an out-of-date view. We can only begin to clean up the mess in our data centers when we have a clear picture of IT assets and dependencies across the organization.



Mapquest for the Data Center

What is needed is an automated way to audit the hardware and software components in the data center, and map business applications onto them – on a continuous basis. Add analytics to help navigate this information, such as reports and dependency graphs – and you have a way to both accelerate and de-risk the planning phase. You also have the means to easily track progress of the optimization program over time.

This is precisely how Tideway can help. Tideway’s dependency mapping product, Tideway Foundation, supports data center projects by providing an automatically generated map of all resources, including server count, operating system count and software count. It automatically detects application dependencies, reducing the impact on the development team and improving planning accuracy. Because it continually updates this map with any changes that occur, plans can be revised to reflect the current state of the data center. Finally, by using change comparison features, it is possible to validate changes to ensure all went as planned.

Because Tideway Foundation is able to achieve 97%+ data quality when recording hardware and software components, many more optimization targets can be identified – translating into very real cost savings. In addition, because all dependencies are understood, the impact of any change is absolutely clear. This takes risk out of change, and unintended impacts on business applications become a thing of the past.

The clear visibility of changes provides a commercial viewpoint to ensure accurate claiming of work carried out, where portions of the project are outsourced. Because, not every data center optimization project is the same, Tideway Foundation comes in different editions to address with different types of optimization project (see sidebar).

Platform edition: This maps the hardware and operating systems in use, making it possible to find old, inefficient hardware and refresh it with resource efficient equipment. This enables rack space reduction and hardware support & maintenance cost reduction projects

Operations edition: This also maps the software on each server, so it now possible to identify hardware and software consolidation opportunities; stranded hardware and software for retirement; virtualization opportunities. In additionally enables software license & maintenance cost reduction

Business edition: Now business applications are mapped to servers, making it possible to understand the business impacts of optimization. This means that whole business applications can be optimized based on business value and application SLA. This enables business application move, refresh and retire projects.

The Tideway Approach

Tideway has extensive experience optimizing global infrastructure projects for Fortune 1000 companies around the world. Its award-winning product, Tideway Foundation, has allowed customers to automate key IT processes by providing visibility of actual configurations and usage which has led to significant savings and data quality improvements over the manual approaches.

Audit

A complete IT inventory – of servers, software and applications, including identifying inter-system and application dependencies – is central to any sustained data center optimization effort. When done manually, this can take person-months to complete at the required granularity and data quality. Data gathered is employed in subsequent planning, costing and risk-mitigation activities, and also serves as the basis of a system of record that will help to manage change during the migration itself.

Project stakeholders are supportive when reassured that rigorous due diligence has taken place with respect to the applications they own, their dependencies to elements of infrastructure and other applications, and that the project team has adopted a complete, professional approach.

Tideway Foundation continuously refreshes a detailed application, software and asset map, including relationships between these components. This builds a clear picture of the components that make up an organization's IT infrastructure, and how those relate to the business.

Analyze

Once a sustainable method of collecting and maintaining accurate inventory and dependency data has been established, organizations can look beyond the data to reach a deeper understanding of how technology underpins its business operations. In this phase, the focus is on addressing questions that relate technology and application assets to their business owners and stakeholders – and mapping out the dependency relationships between them.

Tideway Foundation offers direct support for these activities with its business application modeling capability, dependency graphs and browser-based reporting engine. The information Foundation has gathered automatically is augmented by business contextual knowledge, usually distributed across technology and business silos – and often only kept in peoples' heads. Aggregating all of this in one place can itself be a valuable exercise – and it forms a vital part of the picture.

Design and plan

Armed with a full description of the IT and business asset landscape, the project team can concentrate on finding the optimization opportunities, designing the desired end-state – and planning the process for getting there.

The optimization process is most likely to be executed as a series of discrete steps – ensuring that the overall environment is considered stable and manageable at each checkpoint, rather than as a monolithic and prolonged project which could introduce service level instabilities. The new optimized vision presents an ideal opportunity to implement efficient IT policies to keep the data center in a healthy state on-going.

Designing and planning a data center optimization project is a complex activity with many moving parts. Tideway Foundation's accurate, consistent and sustainable map of the data center infrastructure helps teams to factor and plan their goals accurately. This saves the cost and effort of further manual audits, and indeed, actually makes it possible to evolve plans as the team's insights develop alongside the natural drift of the infrastructure.

Implement

In the implementation phase, Tideway Foundation maintains a precise picture of the state of the IT environment from day to day. This allows project progress to be monitored, provides a continuously updated view of what changed incorrectly in the environment when incidents occur, and assists with necessary re-planning.

If the project has been planned effectively, then the implementation itself will be a controlled, predictable exercise. Tideway Foundation will continuously monitor the data center environment – discovering and tracking:

- Installation and shutdown of servers, including verifying correct configurations
- Installation of operating systems, patch level, kernel level and remote storage
- Installation/decommissioning of software components and business applications

Tideway Foundation's ability to provide extensive up-to-date asset, ownership and dependency data helps to eliminate many of the very costly and time-consuming aspects of keeping an IT inventory current.

Case Study - An Investment Bank optimizing a 1000 server data center

	Before	After
Data quality in system of record	67%	97% (logical) 100% (physical)
Time to perform first audit of data center	250 person-days	2 days
Time for subsequent audits of data center	250 person-days	30 minutes

Benefits	Value
Removal of incremental cost of an audit	\$75,000 saving per one off audit with an incremental audit cost of effectively nil
Identification of additional 30% of server estate that can be virtualized	Additional \$1.8M saving at 8:1 consolidation
Identification of 2.5 - 10% of servers with no function or decommissioned	\$500K - \$2M saving

Summary

Trying to plan data center optimization projects without an accurate picture of the very assets you are trying to optimize is an exercise in futility. The results of this approach are well known – continued data center sprawl, escalating costs, unplanned downtime due to unknown application dependencies. The corresponding lack of data center agility is strangling the business agenda.

By implementing dependency mapping technology, it is possible to get a clear understanding of exactly what is in the data center, and how it relates to business applications. Now planning is based on actuals rather estimates, risk is driven out because dependencies are definitively understood rather than guessed, and the desired cost saving of the optimization exercise can finally be realised in full.