



**APPOSITE**  
TECHNOLOGIES

# Best Practices for Data Center Relocation:

## Manage Expectations and Get It Right the First Time

### INTRODUCTION: THE PROMISE

Packing up and moving the resources in your data center means uprooting the digital brain-trust and nervous system of the business. And, it costs a lot. According to Info-Tech Research Group, "The average data center relocation costs \$120,000 or \$10,000 per rack."

So why do we do it? The reasons are compelling and diverse:

- Mergers & acquisitions
- Business continuity
- Cyber security
- Virtualization
- Information Technology (IT) cost reduction and operational efficiencies

With the potential upside both lasting and huge, data center consolidation continues as new trends take hold. An article by Data Center Knowledge, a division of the Informa, cited several leading experts offering their take on the future.

According to the article, IDC saw the number of data centers worldwide continuing to climb until 2015, when the number reached 8.55 million. The worldwide number would decline to 8.4 million in 2017, IDC predicted last year, and again to 7.2 million by 2021.

The same article quotes Mike Leone, senior analyst at Enterprise Strategy Group, explaining that current efforts focus on "not only reducing the hardware footprint, but accelerating virtualization adoption to deliver a cloud-like experience on-premises" and making "it easier to virtualize, easier to manage, easier to scale, and do it all cost-effectively." This goal signals a shift from consolidating lots of smaller data centers into large regional facilities to that of scaling down on real estate and power costs through virtualization.

Whatever the goal, delivering on the promise of data center relocation and consolidation entails moving resources from one place or platform to another. This process invariably involves significant risk, change, and apprehension among

users as well as IT teams, and can impact applications as well as servers, in critical ways:

- Local users may suddenly become remote users
- The user experience may suffer as application response times increase due to longer round-trips to and from servers
- Server performance, and in turn, application scalability, may be degraded as servers handle more remote transactions that take longer to complete
- Migration to the cloud introduces new challenges for visibility, security, and compliance

In fact, 20-40% of applications will fail to meet service level objectives during or after a data center relocation, some with severe issues that render the application unusable.

Addressing service level issues upon migration may mean additional changes to server systems and even the re-coding of applications – both of which add delay, risk, cost, and frustration. But it doesn't have to be that way.

## WHAT GIVES RISE TO RISK?

You budget. You plan. You execute—and watch for the unknowns that can come to light at any time. The trick to minimizing the risk impact of unknowns is to eliminate as many as possible by modeling and thoroughly assessing the “knowns.” Priority areas of risk avoidance include:

### Application performance/transaction times

Moving servers to a new location obviously changes the means by which applications get delivered over the network. Users previously considered local to directories, domain servers and desktop services now become remote. With servers now being accessed from a distance, transaction response times increase for everything from booting up and logging in to sharing and storing large files or repositories of data.

Network and geographic latency created by the distance between clients and servers gets compounded by higher numbers of users transacting with servers remotely via wide area network (WAN) connections. To users, the network latency manifests as application latency which manifests for IT as a rise in complaints and trouble tickets.

## Mitigate Risk, Meet Expectations

The key to side-stepping the pitfalls and realizing the promise of data center relocation is to eliminate surprises and address potential issues beforehand. Doing so means:

- Baselining application performance prior to the move
- Modeling and measuring performance of the new infrastructure before physically relocating
- Continuing to validate service levels post-migration

Ongoing testing includes measuring network latency and other impairments and their impact on applications, business processes, and users. The insight derived from these activities also serves to fuel better plans for handling data backup, replication, security, and data compliance.

In this paper, we will focus squarely on best practices for predicting and maximizing application performance throughout data center relocation to meet user expectations and service level requirements out of the gate. To do so, a deeper understanding of the risks proves useful.

Each application will be impacted differently depending on their unique sensitivity to latency based on things like:

- **Application chattiness:** How many messages must be exchanged between the client and the server for each transaction?
- **Blocking:** How many turns can be processed in parallel?
- **Protocol configuration** for both clients and servers (buffering, authentication, etc.)

Careful planning and modeling of individual application performance in the new target environment should be conducted in the lab to prevent noticeable degradation of the user experience.

### Server performance

During data center relocation or consolidation, servers supporting a given number of local users may suddenly become accessible to many more remote users conducting slower transactions due to increased distance and network latency. Web-based applications in particular may result in exponentially more users performing low-bandwidth but high-latency processes.

These higher volumes of transactions queuing up and taking longer to complete hinder the performance and scalability of servers themselves. Because servers allocate and lock in resources for the duration of a transaction, added latency causes transactions to complete more slowly, tying up more server resources for longer periods of time. Higher volumes of in-flight transactions can eventually exceed a server's capacity resulting in critical transactions or data being dropped.

Simply throwing more CPU power at the problem, however, does little good. Subtle changes to processes scheduling may be required, and it may even make sense to explore the option of offloading processing-intensive functions such as decryption/encryption to dedicated appliances.

Once again, it is critical to model all of this in the lab beforehand because, once the move takes place, it becomes harder to replicate and troubleshoot the specific conditions causing problems. Server capacity planning teams must model and test performance ahead of time by simulating a realistic mix of local and remote users running real applications. QA teams must also consider latencies between users and servers more carefully and in a new light to support optimal budgeting and resource allocation.

## Finger-pointing

While the classic reaction to poor application performance is "blame the network," application architects, developers, and quality assurance (QA) teams must also understand and design higher tolerance for latency into applications. Issues linked to subtle yet complex interactions between applications, networks and infrastructure cannot be resolved by working in silos. Unprecedented collaboration must happen between application owners, systems managers and network architects.

Careful planning and detailed "before" and "after" insight into performance equips project managers to join forces to meet corporate goals for performance, time to value, and return on investment (ROI) through data center relocation.

## Business continuity

The larger the center, the longer it might take to safely and systematically move it. Centers containing hundreds or thousands of servers could take weeks, months, or longer with systems being

maintained in two places for an extended period of time.

This interim stage means that relocation managers must factor in challenges around:

- Inter-dependencies on back-end systems temporarily located in different places
- Maintaining (or splitting up) server groups
- Replication and backup

The team must contend with issues arising from temporary latencies between critical back-end systems as well as those created by clients and servers being in different places. Both scenarios can wreak havoc on application performance, especially when server systems are not designed to accommodate latency.

Server groups and dependencies must be fully considered in planning data center relocations to avoid costly, highly visible impairments of business processes.

## Service-level expectations

The very idea of change often elicits a negative reaction, at least at first. Who wants to risk having things break, leak, or take longer, even temporarily, in hopes of realizing gains that seem to mainly benefit IT, or "corporate"? Who wants to go first and be the guinea pig for the inevitable "fine tuning"?

Again, the surest way to assuage these very legitimate fears is with very clear data:

- How exactly will applications perform during and after the move?
- How long will moving take? How much will users be inconvenienced during the transition and how good will IT support be throughout?
- What might the downtime cost various business groups?

Business concerns such as these can cause rollbacks, delays, and a predisposition to complain during and after the process. Being able to provide clear, reproducible data showing exactly what the impact on users and applications will be beforehand goes a long way in alleviating concerns, as well as meeting expectations and demonstrating value to stakeholders. Even when performance trade-offs are necessary, the ability to proactively manage expectations will help prepare users and may inspire better methods.

## SO, HOW DO YOU KEEP YOUR PROMISE?

At any given time, users are accessing resources from regional and branch offices, from home and while on the road, connecting into your network through a mix of private lines, the public Internet, virtual private networks (VPNs), and clouds. Given the many variables, applications are rarely as responsive over the WAN as they are at headquarters with the impact ranging from “slight frustration” to “wholly unusable.”

Bandwidth constraints and bottlenecks, latency, jitter, packet loss and other impairments can all wreak havoc on application throughput and responsiveness. Traditional testing on the local network fails to identify such issues that may impact the end-user experience and, worse yet, simply increasing WAN bandwidth may do little to improve things. Even WAN acceleration may not be enough.

Enter WAN emulation. This critical best practice in rolling out or migrating networks and applications makes it simple and affordable to test performance in the lab under real-world conditions. WAN emulation streamlines data center relocation, virtualization and other migrations by answering crucial questions such as:

- How responsive will the database, ERP, inventory, and order fulfillment systems be to users in branch offices?
- Will Sales be able to use the CRM system from the road?
- How will VoIP sound halfway around the world?
- Are cloud-based, wireless or satellite networks viable alternatives for remote users?
- How much bandwidth is really needed to keep users working productively?

Without WAN emulation, network managers traditionally tried two approaches to answering these questions: testing applications on the local network and limited testing on the production network. Neither suffices for data center relocation.

For one thing, the chatty, transactional nature of HTTP, CIFS, and other client/server communications may be severely impacted by latency and other network impairments that cannot be measured simply by testing on the local network. Shipping equipment and forth to test on the production network riskier, not to mention more tedious and

costly approach. Testing on the live network also may be limited to off-peak hours between a few convenient sites, and often misses problems that occur under more typical or challenging circumstances.

A powerful and far simpler approach, WAN emulation lets IT and networking teams see what users will see before they see it.

### Emulating the network in the lab

Until recently, emulating WAN links required test specialists with expensive, complex telecommunications hardware unfamiliar to most IT groups. A next-generation solution—Netropy WAN emulators from Apposite Technologies—introduces a new approach to network simulation architected to bridge the gap between network and IT operations.

By adapting commodity hardware with high-performance packet processing algorithms, Netropy emulators feature the precision of a high-end test tool in a cost-effective, easy-to-set-up-and-use appliance delivering valuable insight throughout the application lifecycle:

**Network design:** Rollouts, adding links, and increasing bandwidth all impact application performance. Private networks, internet VPNs, satellite links, and wireless connectivity each present trade-offs in cost, performance, and convenience that can be thoroughly modeled and better understood by simulating expected usage and evaluating “pros” and “cons.”

During data center relocation planning, emulation should be used to assess how many concurrent users a server will be able to support post-migration without risking dropped sessions or introducing intolerable transaction latency. This means simulating varying combinations of local and remote users connecting over different types of networks.

**Application validation & benchmarking:** No two applications are exactly the same. Large file transfers may be sensitive to link bandwidth, latency and loss while database and web-based applications can be highly sensitive to congestion as well as latency. One issue is best understood in terms of throughput rates while the other centers on responsiveness. Still other real-time applications such as voice and video can be impacted by jitter, congestion, and other transient effects as well as by latency and bandwidth constraints.

Emulating conditions and evaluating performance under a variety of real-world conditions helps to refine plans and benchmark the offerings of competing vendors.

**Application Optimization:** Improving performance during product updates and by tuning internal parameters. The Netropy WAN emulator provides network managers with a convenient test bed to determine the impact of modifications to any application.

**Troubleshooting:** Seeing what users see by replicating the network to help reproduce problems in the lab. Once the problem is identified, fixes or workarounds can then be validated on the test network.

**Moving to the cloud:** Bandwidth constraints and latency can heavily impact the time it takes to complete migration. To streamline the transition, organizations must ensure there is enough capacity to handle the new cloud configuration. Network emulators help evaluate how network constraints such as bandwidth limitations and latency will impact a large data transfer without the risk of losing data.

### The Netropy WAN Emulator Simulates:

- ✓ Bandwidth
- ✓ Latency and Jitter
- ✓ Packet Loss
- ✓ Background traffic

## A STEP-BY-STEP GUIDE TO SMOOTHER RELOCATION

Following the steps below in the order shown will help ease transitions and reduce risk, cost and frustration by introducing predictability and insight into the process:

**1. Identify applications that will experience performance degradation.** Simulate the new data center environment by programming network conditions between the user's desktop and application servers post-migration.

**2. Investigate issues.** Isolate the specific conditions and bottlenecks impacting each application.

**3. Remediate performance problems and validate solutions in a proof of concept lab.** Use remediation labs to test the effectiveness of proposed solutions such as code fixes and WAN acceleration to measure performance gains.

**4. Assess impact on interdependent back-end servers** Understand and proactively address latencies that may be created between servers during transitional stages of the move using network emulation to measure the impact of these server-to-server latencies.

**5. Manage user expectations.** Let users experience post-move application performance in advance and take steps to mitigate or address their input and concerns.

Use WAN emulation throughout the process to over-deliver and speed time-to-benefit by turning unknowns into knowns that can be explained and managed.

## CONCLUSION

Relocation project managers can take the mystery, risk and apprehension out of data center relocation using proven best practices for planning, design and deployment. Simulating each stage of the move ensures the business is properly supported throughout. Continuing to emulate WAN conditions after migration improves operations and speeds troubleshooting while reducing risk, effort and cost.

WAN emulation bridges the expertise gaps between IT, network and security teams with a practical tool any engineer or administrator can use to level-set expectations and collaborate on new processes and best practices. Perhaps more important, it provides a framework for minimizing any real or perceived unwanted impact on users, and for demonstrating that methodologies and results are delivering on the promise of relocation.

In other words, WAN emulation plays a critical role in amassing the insight needed to get data center relocation right the first time, and continue adding value over time.

## About Apposite Technologies

Apposite Technologies makes WAN emulation easy by offering professional-quality network emulation tools at affordable prices. Apposite's award-winning Netropy and Linktropy WAN emulation appliances simulate bandwidth, latency, loss, congestion, and other network impairments with fine-grained precision to provide accurate simulations of any type of wide-area network. Netropy and Linktropy WAN emulators are widely deployed by leading enterprises, application and equipment developers, telecoms carriers, and government and military organizations around the world. Apposite Technologies – WAN Emulation Made Easy



**APPOSITE**  
— TECHNOLOGIES

**17835 Ventura Blvd Suite 211, Los Angeles, CA 91316 USA**

**TEL: 1.310.477.9955 | [info@apposite-tech.com](mailto:info@apposite-tech.com)**

**[www.apposite-tech.com](http://www.apposite-tech.com)**

Copyright ©2019 Apposite Technologies LLC.

All rights reserved.

Apposite, Linktropy and Netropy are registered trademarks of Apposite Technologies.

The Apposite logo and "WAN emulation made easy" are trademarks of Apposite Technologies.

P/N: DOC-DSNNE-071