THE CLIPPER GROUP
Captain's Log™
Navigating Information Technology Horizons℠

Published Since 2001 | Report #TCG2003019 | May 8, 2003

On-Off Capacity on Demand —
Does This Make Sense for Your IT Infrastructure?

Analyst: Mike Kahn

Having the sufficient IT resources to meet planned and unplanned enterprise requirements is like walking on a tightrope – lean a little too much to the left and you fail because you have under provisioned; lean a little too far to the right and you fail because you have overprovisioned. **What you need is a way increase the probability that you can traverse the wire without falling, without the cost of extensive overprovisioning.**

Good IT managers understand the expected demands of an application throughout a business cycle. They know when the spikes for resources are likely to occur, be it end-of-month processing or seasonal promotions. What they may not know, for some spikes in demand, is how big they will be. For example, if a seasonal promotion is driven by advertising, how successful will it be in stimulating inquiries and orders? Then there are the unexpected demands, driven by external forces beyond the enterprise's control. For example, what if there is a tornado tearing through the Midwest and the enterprise has to process 100 times the normal weekly insurance claims?

What these have in common is the need for sufficient *reserve power* to meet the spikes in demand. It's easy to solve this problem on an unlimited budget, a luxury that few have during these economically stagnating times. **The challenge is to be reasonably prepared at a reasonable cost.**

If you were a general contractor and you needed a few extra trucks (or generators, aerial buckets, etc.) for an especially large project, would you go out and buy them? You might think about it, in terms of future needs to determine whether you would be getting a good ROI, but most likely you would just rent what you needed for the short term, and let someone else worry about recovering the investment on the asset.

This is all about renting to meet your occasional incremental requirements. To do this successfully, you need to know:

- Your baseline workload(s)
- Your expected (cyclical) spikes
- Your greatest fear, in terms of workload demand
- Your ability to shed less critical workload, when workload spikes
- The risk that your enterprise faces, if it cannot meet demand

With this information in hand, then you need to ask the killer question: ***How can I add to my capacity on demand - with the assurance that the spare capacity will be available - without paying through the nose?***

The Clipper Group, Inc. - Technology Acquisition Consultants ◆ Strategic Advisors
888 Worcester Street ◆ Suite 90 ◆ Wellesley, Massachusetts 02482 ◆ U.S.A. ◆ 781-235-0085 ◆ 781-235-5454 FAX
Visit Clipper at www.clipper.com ◆ Send comments to editor@clipper.com

As mentioned earlier, you could choose to buy enough infrastructure to meet any possible need, and let it sit idle waiting for the need to arise. This is what many brokerage firms have done, because the risk is too high if they are unable to handle the next really-big trading day. You could overprovision a little, and plan to handle the spike by cutting back other work, but that is getting harder to do in a 7x24 global economy. Or you could take advantage of vendor's options to upgrade capacity on demand, which would allow you to turn on extra processors that were sitting as spares in your SMP server(s). Capacity upgrade on demand (CUoD) is attractive, because you can turn it on quickly[1], but in almost all cases once it is turned on, it stays on, with likely impact on software charges from that point forward.

Think back to the general contractor example. If the foreman knew that the local rental store would always have what he wanted when he wanted it, that the time to get it to the site was acceptable, and that the price was reasonable, then he would spot rent the needed items. Assuming that the price could be determined, in advance, to be reasonable, and that time to delivery was not an issue, then the big question is availability. If it isn't available when you need it – and you don't really know when you are going to need it – renting isn't going to reduce the risk, which is the primary requirement for meeting the spike in demand.

So what you want is the ability to turn on the extra resources when you decide that they are needed, and to turn them off when the spike is over; all for a reasonable fee. This is called *On-Off Capacity on Demand*. It is another option to consider in provisioning and managing an enterprise's IT infrastructure. While there is technology involved to quickly add capacity without shutting down the larger system, **On-Off Capacity on Demand is more about financing and risk abatement**.

- **You need to assess how long your spikes will run** (hours/days/weeks), before returning below your baseline capacity. If the source of the spike is normal, on-going application or business growth, then it will be more economical to buy the additional capacity than to rent it.
- **You need to estimate the range and probabilities of the size of your spike**, to determine the quantity of resources that you might want to rent.
- **You need to determine if your vendor offers On-Off Capacity on Demand** for the resources that you think you might need to meet your spikes, and
- **You need to carefully review the terms and conditions that will apply**.

The thing that you want to avoid is using the rental option for too long, i.e., to have paid for it in rental charges without the benefit of owning it.

On-Off Capacity on Demand is a further example in the trend toward partnering with your suppliers, to the benefit of both.

- **Is it available broadly?** Not yet. This is a new concept.
- **Does it apply to all IT resources?** No, unless you outsource the entire data center. Look for it in SMP servers where CUoD already exists (and you are likely able to turn off a resource – think *processor* or *memory* – when the spike is over (this may not be applicable to storage, since capacity almost never goes down).
- **Is it worth considering?** Yes, as a financially-leveraged risk abatement vehicle.
- **Is it easy to assess?** You may need some help.

With budgets tightening and demand growing, On-Off Capacity on Demand may give you some breathing room. Put this one near the top of your to-do list.

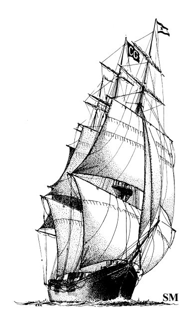

[1] For servers, exactly how quickly varies by operating system, server and vendor's CUoD offering.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Mike Kahn **is Managing Director and a cofounder of The Clipper Group.** Mr. Kahn is a thirty-year veteran of the computer industry. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation — Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; President of Power Factor Corporation, a Boston-based electronics firm; at Honeywell Bull — Director of International Marketing and Support; at Honeywell Information Systems — Director of Marketing and Director of Strategy, Technology and Research; with Arthur D. Little, Inc. — a consultant specializing in database management systems and information resource management; and, for Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan where he earned B.S.E. and M.S.E. degrees in industrial engineering.

- ***Reach Mike Kahn via e-mail at Mike.Kahn@clipper.com*** **or via phone at** ***781-235-0085 Ext. 21. (Please dial "1-21" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.