

For: CIOs

# Your BT Agenda Demands A Hybrid Cloud Storage Strategy

by Henry Baltazar, January 8, 2015

## KEY TAKEAWAYS

### **Current Cloud Storage Deployments Are Limited**

The majority of cloud storage deployments were built either for cloud-native applications or to handle unstructured data for low-cost archiving and backup retention. Classic physics and other hurdles complicate the cloud storage puzzle for typical distributed enterprise-grade business applications.

### **Workload Migration Is The Next Frontier For Hybrid Cloud**

Leveraging dedicated cloud connectivity at colocation sites and cloud gateway technologies, hybrid cloud storage will facilitate workload mobility both from on-premises to clouds and between cloud providers.

### **Hybrid Clouds Will Enhance Collaboration For Highly Distributed Organizations**

Global file locking and caching techniques will allow clients to expand remote collaboration to more intensive workloads such as engineering and CAD while compensating for slow and inconsistent links at remote worksites.



## Your BT Agenda Demands A Hybrid Cloud Storage Strategy

How You Handle Storage Will Heavily Influence Your Cloud Success

by [Henry Baltazar](#)

with [Glenn O'Donnell](#), [Dave Bartoletti](#) and Andrew Hewitt

### WHY READ THIS REPORT

Hybrid clouds are the future state of compute and storage for enterprises in the age of the customer. In recent years, hybrid cloud storage has evolved and become a deployment methodology that can now blend on-premises storage infrastructure with public and hosted cloud services. The road to hybrid clouds will not be a short or easy one for many organizations, and we expect that many will have to make detours along the way to take advantage of the rapid innovation of hyperscale cloud storage environments. The workload and data mobility that can be created with hybrid cloud storage strategies will allow infrastructure and operations (I&O) professionals to hedge their cloud bets and switch to new providers when terms or capabilities no longer meet expectations or when cloud storage providers offer a new “can’t miss” service.

### Table Of Contents

- 2 **Current Cloud Storage Environments Are Limited**
  - Location Matters: Drawbacks To A Mass Cloud Storage Migration
  - Evolving Colocation And Caching Strategies Pave The Way To Workload Migration
- 4 **Raise Infrastructure To The Next Level With Hybrid Cloud Storage**
  - Shared Storage Unlocks the Full Potential Of Hybrid Clouds
  - Accelerate Collaboration With Caching And Global File Locking
- RECOMMENDATIONS**
- 8 **Embrace Hybrid Architectures To Prepare For Your Exodus To Cloud**
- WHAT IT MEANS**
- 9 **Cloud Forces A New Architectural Approach To Everything**

### Notes & Resources

Forrester leveraged input from inquiries and briefings, in addition to past Forrester research, in the formation of this research report.

### Related Research Documents

[Velocity: The Only Storage Issue That Matters In The Age Of The Customer](#)  
November 26, 2014

[Market Overview: Software-Only Storage](#)  
August 1, 2014

[What The Evolution Of Cloud Storage Means For I&O](#)  
February 14, 2014



## CURRENT CLOUD STORAGE ENVIRONMENTS ARE LIMITED

Make no mistake. I&O professionals cannot build a comprehensive enterprise cloud compute strategy without first building a foundation based on hybrid cloud storage technologies. While the elasticity and power of cloud compute and networking resources are game-changing capabilities, these services will have limited use beyond test and development scenarios if you cannot leverage your company's or partner's data or if you cannot migrate existing workloads.

Thus far, enterprises have treated cloud storage resources as complementary storage silos to support next-generation analytics and systems of engagement cloud applications, and to offload menial unstructured storage burdens. But cloud storage has much more to offer. Next-generation hybrid cloud storage technologies are already available and ready to bridge the gap between your on-premises applications and resources and cloud storage services by simplifying workload mobility and cloud federation between service providers.

The footprint of enterprise data resident in cloud storage services has grown dramatically with the popularity of file sync and share and object storage services such as Google Cloud Storage and Amazon S3. Still, physics and WAN latency have limited the use cases for cloud storage. Conventional on-premises enterprise applications — such as databases — are designed with the assumption of proximity between compute and storage resources, and expect latencies on the order of milliseconds that are not possible over WAN links. As a result, cloud storage has been limited to two basic scenarios:

- **Secondary storage for unstructured data.** “Cheap and deep” storage was an early driver for the adoption of cloud storage, but in recent years, organizations have also leveraged the networking resources of hyperscale object storage services such as AWS's S3 to facilitate the delivery of storage services to mobile devices and users. WAN latencies on the scale of hundreds of milliseconds are tolerable for the retrieval of documents and other forms of unstructured data.
- **Cloud storage for cloud compute.** The addition of all-flash storage service offerings such as Amazon's Provisioned IOPS for EBS (Elastic Block Store) or Google's SSD Persistent Disks provide high-performance block-level storage for running cloud compute applications. While this capability has made it possible to create transactional applications in cloud environments, these block storage services can only be used by cloud compute instances residing in the same cloud and availability zone.

## Location Matters: Drawbacks To A Mass Cloud Storage Migration

The rapidly falling price of cloud storage services is pushing organizations to migrate away from expensive conventional storage systems and toward cloud storage services. While cloud storage should absolutely be a part of your future storage infrastructure, there are a few limitations that should be kept in mind, including:

- **Long migrations and lock-in.** The untimely collapse of Nirvanix forced its clients to hastily move away from its cloud storage services, shocked early adopters, and clearly highlighted the need for cloud federation and data migration services.<sup>1</sup> Long-distance data migration continues to be a struggle for enterprises, specifically those with large workloads in the scale of hundreds of terabytes to petabytes, and this is exacerbated by the inconsistent network connectivity commonly found in remote sites. Data transfer across IP networks can take a significant amount of time when data payloads go beyond 100 terabytes (see Figure 1). The cost of outbound bandwidth and time associated with migrations are deterrents that make cloud migrations difficult and effectively lock in clients to a provider.
- **Security and compliance.** Cloud security is still the No. 1 impediment to full-scale cloud service adoption, though the steady stream of improvements cloud providers are making to their environments are countering this objection.<sup>2</sup> Ultimately, enterprise storage teams are hesitant about handing over control of data to cloud providers and use security and compliance concerns to slow down cloud deployments.

**Figure 1** Big Data Migrations Take Time And Bandwidth

Connectivity	Time to transfer 1 TB	Time to transfer 100 TB	Time to transfer 1 PB
10 Gbps link	12.42 minutes	20.7 hours	8.6 days
1 Gbps link	124.18 minutes	8.6 days	86 days
100 Mbps link	21 hours	88.3 days	883 days

\*Times are based on theoretical rates. Actual transfer times will be two to three times longer.

117944

Source: Forrester Research, Inc. Unauthorized reproduction or distribution prohibited.

### Evolving Colocation And Caching Strategies Pave The Way To Workload Migration

While the physics limitations of rapid data movement cannot change, progressive enterprises and service providers have come up with clever new strategies to put compute and storage resources closer together, including:

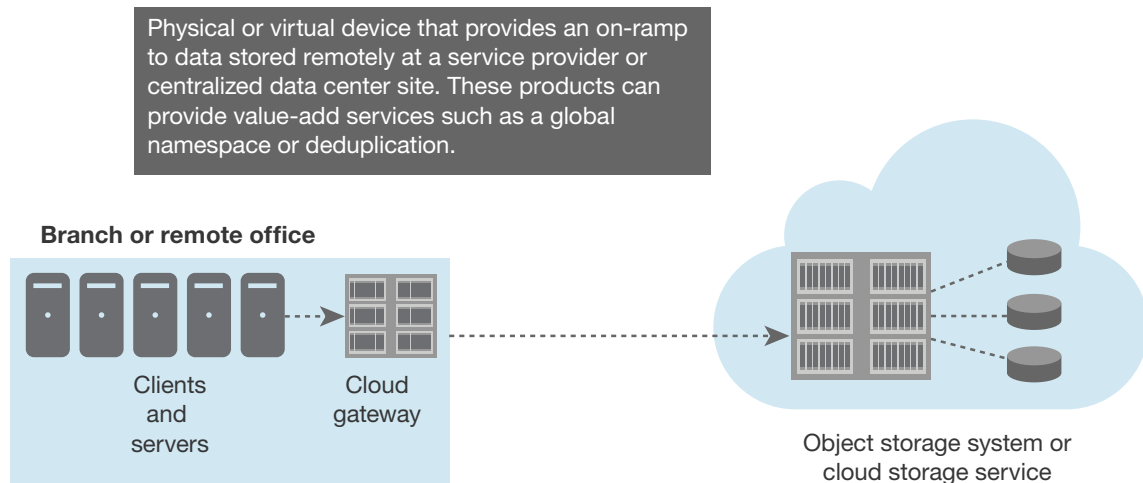
- **Accelerated network connectivity with colocation.** A handful of colocation and data center providers, such as Equinix, CoreSite, and TelecityGroup, have facilities that are either physically close to a cloud service provider or house an instance within their data center. This proximity allows customers to bypass the Internet and connect directly to their cloud service providers (often for a flat monthly rate, via fiber-cross-connections at port speeds of 1 gigabit per second or 10 gigabits per second), which can provide as low as sub-millisecond latency since such

networks bypass the Internet.<sup>3</sup> Amazon's Direct Connect emerged in 2011 as a means of accelerating data movement, and its partner Equinix claims that Direct Connect can transfer files to Amazon 138% faster than public Internet connections.<sup>4</sup>

- **Cloud gateway/controllers.** These devices have been on the market for a few years but so far are not widely deployed. Cloud gateways are virtual or physical devices that sit on-premises at a client's site and funnel data off to cloud storage services or object storage systems (see Figure 2). These systems provide local caching to minimize the amount of data that must be fetched from cloud storage over Internet links. Forrester predicts a surge of activity in this space in 2015 with the accelerating growth of current players Ctera Networks, Nasuni, and Panzura, and with the major storage players recently making acquisitions in this space, such as NetApp's acquisition of Riverbed Technology's SteelStore gateway and EMC's acquisition of TwinStrata.

While cloud gateways and network-accelerated colocation have been around for a few years, they are just now getting closer to mainstream adoption for enterprises. These technologies will be key enablers for the next wave of hybrid cloud deployments.

**Figure 2** Cloud Gateways And Controllers Funnel Data To Clouds



## RAISE INFRASTRUCTURE TO THE NEXT LEVEL WITH HYBRID CLOUD STORAGE

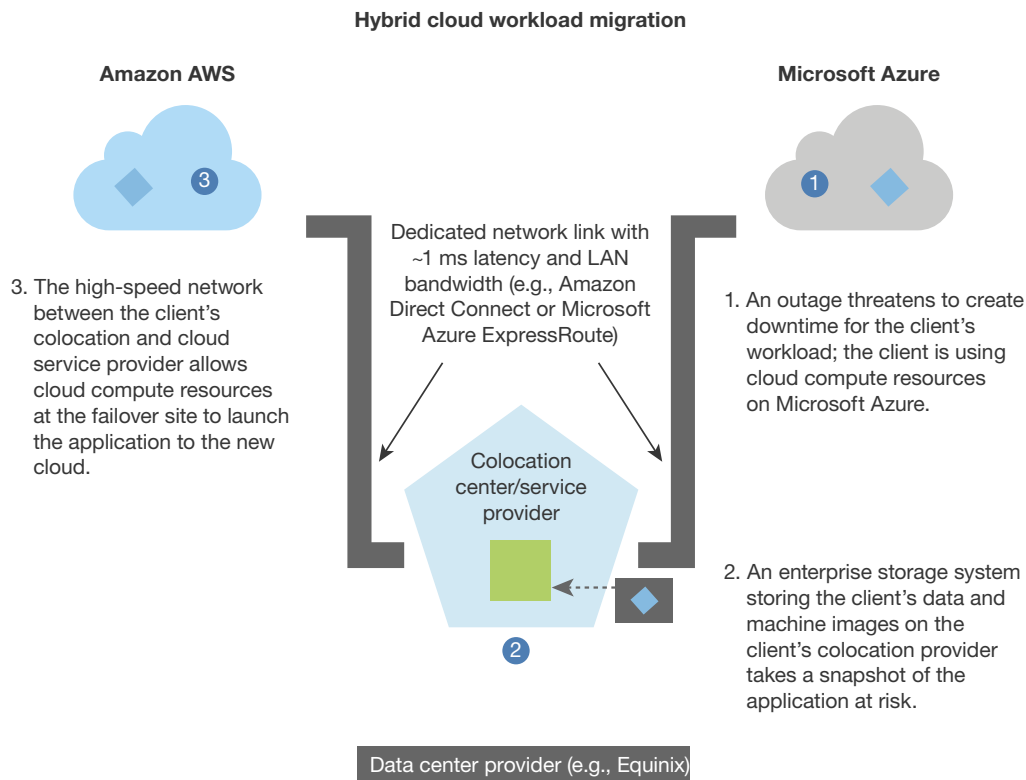
Improved data sharing, mobility, and elasticity are the fundamental benefits for hybrid cloud storage environments. By leveraging the accelerated networking and cloud gateway caching technologies discussed in the prior section, organizations will be able to expand the flexibility of their cloud deployments.

## Shared Storage Unlocks The Full Potential Of Hybrid Clouds

Shared storage volumes are a fundamental resource in enterprise and midrange environments and are necessary for running clustered applications such as Oracle Real Application Clusters (RAC). They are not a standard service with many cloud providers, however. Popular block storage services such as Amazon's EBS can only mount an EBS volume to a single compute instance. Enterprise storage vendors such as NetApp (with its NetApp Private Storage for Cloud solution family) and Zadara Storage are teaming up with colocation service providers to facilitate instant data sharing across clouds without data migration. To take your cloud application storage to the next level, consider the following steps:

- **Create a stepping stone to clouds.** The shared storage cloud deployment allows customers to apply familiar enterprise storage feature sets such as volume sharing, optimization (compression/deduplication), and application integration to their cloud workloads. The costs of these deployments will be substantial since they will entail storage system, physical rack space, and networking costs. On the positive side, for bursty and seasonal workloads, organizations should see a substantial cost savings for compute since they will be taking advantage of the elasticity of cloud services. For risk-averse enterprises, the cost of the deployment can be justified by the ability to leverage familiar data protection and storage management capabilities while learning how to optimize workloads in cloud environments. Also, switching cloud providers is likely at some point in the future, so consider colocation networks to make the switch easier.
- **Add shared storage to accelerate cloud performance.** By placing enterprise storage systems in colocation facilities with accelerated network connectivity — such as AWS Direct Connect or Microsoft's Azure ExpressRoute — machine instances in clouds will be able to access data at the colocation site without forcing a data migration. The initial instances of these deployments were created to provide a higher performance alternative to cloud storage services like Amazon EBS.
- **Facilitate failover across clouds.** The ability to failover workloads across clouds will become more common in 2015. This past year, initial demonstrations of this deployment allowed workloads to failover from one cloud to another without data movement (see Figure 3). The principles behind how this layout works are similar to how VMware's vMotion and Microsoft's Virtual Machine Live Migration migrate VMs between physical servers by changing access rights and presenting the data set to the replacement server. In the colocation site, the high-speed connectivity options available provide low-latency access and performance on 10-gigabit Ethernet links comparable to conventional Fibre Channel SAN links. This makes workload migration possible across clouds.
- **Tap into hybrid cloud compute.** The data-sharing capability of these hybrid storage architectures will facilitate hybrid compute. Given the rapid evolution of services in cloud environments, the ability to present shared storage to new cloud compute or analytics services will give organizations the ability to tap into these services without migrating their entire data set to a particular cloud.

**Figure 3** Leverage Hybrid Cloud Storage To Migrate Workloads Without Moving Data



117944

Source: Forrester Research, Inc. Unauthorized reproduction or distribution prohibited.

### Accelerate Collaboration With Caching And Global File Locking

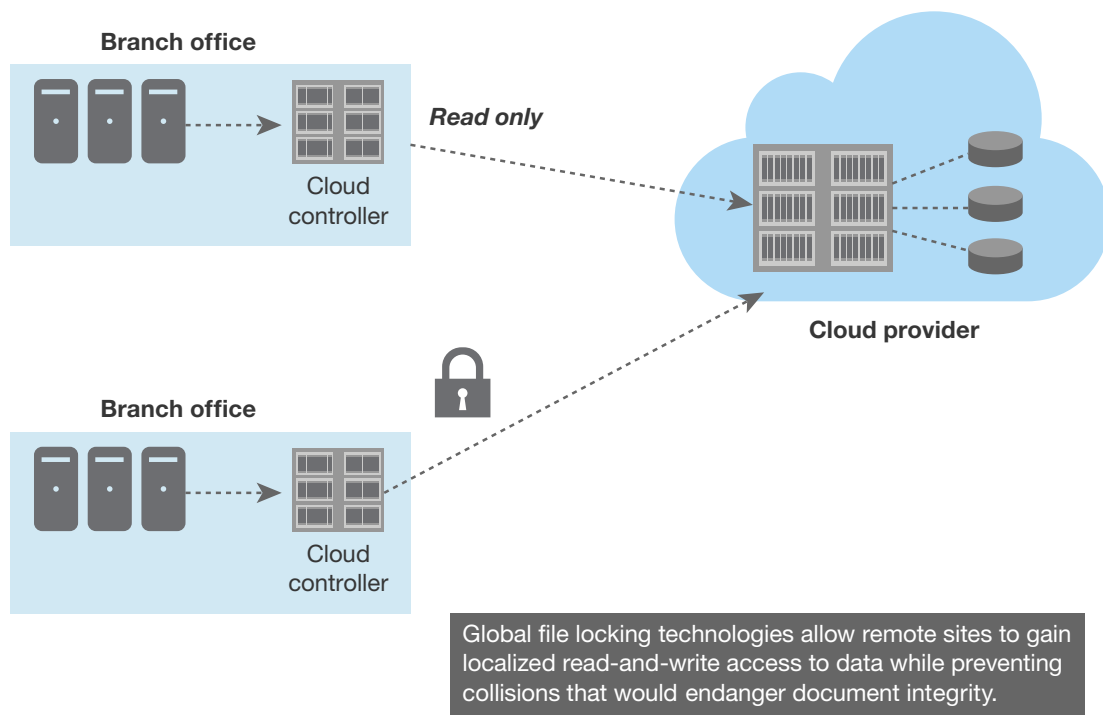
Sharing large files at remote sites is a difficult problem that many sites still struggle with. This problem only gets worse when users at remote sites need to simultaneously edit the same file set. Global file locking is a new feature that cloud controller vendors have added to combat this issue, which has been an acute problem for clients with ACE (architectural, engineering, and construction) projects. Technologies such as intelligent local caching and file locking are powerful because they can provide:

- **Latency reduction.** Latency and packet drops cripple data access for remote sites. For example, in the architecture and construction space, C&S Companies, using Panzura controllers, reduced its file open access time from 22 minutes to 8 seconds by leveraging the caching and global file locking capabilities of the gateways (see Figure 4).<sup>5</sup> C&S used the deployment with key ACE applications such as Revit, AutoCAD Civil 3D, and Bentley MicroStation, which are slowed

down by the thousands of roundtrip transactions required to open, lock, and close files. With files that are highly active, global file locking can prevent accidental file corruptions and collisions since only one user will be able to have write access at any given time.

- **Global data access at remote offices.** JWT, a WPP company, turned to Ctera Networks cloud gateways to provide data access when political instability at its offices in the Middle East forced its office workers to relocate. The firm centralized the data storage for its eight Middle Eastern offices and 800 employees in Amazon AWS S3 cloud storage and leveraged the security capabilities of the cloud to allow the country manager of each office to control its data sets to meet the company's compliance guidelines. The caching provided by the devices enabled data access despite expensive and inconsistent broadband lines that were available to the offices. It also facilitated data capture at the remote sites, which was invaluable since there are high employee turnover rates at these sites. JWT claims it was able to save 63% compared with its previous tape-based backup deployed at the branch offices.

**Figure 4** Global File Locking Can Enhance Collaboration On Cloud Storage





---

## RECOMMENDATIONS

### EMBRACE HYBRID ARCHITECTURES TO PREPARE FOR YOUR EXODUS TO CLOUD

Clouds and the workloads they house are getting larger and more business critical on a daily basis. Beyond raw scale, you must leverage new strategies and technologies and super charge your hybrid cloud storage strategy to keep pace with innovation. To help make this evolution to hybrid cloud storage a smooth one, forge a strategy that:

- **Facilitates cloud mobility.** The cloud computing market is evolving rapidly, with new services and price wars increasing competition. Any supporting hybrid cloud storage strategy should leverage resources such as AWS Direct Connect, Azure ExpressRoute, NetApp Private Storage for Cloud, and Zadara Storage to make data available to multiple cloud services.
- **Prioritizes resiliency.** Outages of cloud services continue to occur, and you should not expect them to end any time soon. As mission-critical and business-critical workloads make their way into the cloud, your teams should leverage multiple availability zones and colocation sites to ensure that workloads can continue to function in the event of a failure.
- **Is mindful of bandwidth costs.** Premium network charges for DirectConnect and ExpressRoute will save money over the long run if large data sets must be moved to or between clouds. Migrations should be cost justified and planned accordingly because these operations could take days or weeks to complete. While inbound data charges have been eliminated over the past few years, the fees for outbound network traffic can become significant.
- **Classifies and optimizes workloads.** While new deployments such as NetApp Private Storage are bringing powerful data sharing capabilities to the mix, these deployments can become costly, as they require a number of expenses, such as colocation space, networking charges, and the deployment of a NetApp storage system. Prioritize the workloads that require hybrid cloud computing capabilities and failover for the shared storage deployments at colocation sites and leave less intensive workloads in the cloud to control costs.
- **Takes security into account.** The security of data in cloud environments is a shared responsibility between clients and their service providers. The providers can only be held responsible for securing the basic infrastructure, which they do extremely well. You are still liable for your application, governance (i.e., determining who has the authority and responsibility to make changes and how to make them), life-cycle management, and key metadata for configurations and dependencies. As your organization's use of cloud resources expands across geographies and service providers, security must meet the same standards of your primary data centers.

- **Tests failover and workload migration operations often.** The pay-as-you-go pricing of cloud computing will allow you and your team to test workload migration without a costly hardware expense. Test early and often to ensure that your disaster-recovery-as-a-service (DRaaS) and cloud deployments are configured correctly and can handle the workload when a migration or failover is necessary.

---

#### WHAT IT MEANS

### CLOUD FORCES A NEW ARCHITECTURAL APPROACH TO EVERYTHING

Cloud services are inevitable because they bring profound benefits. The classic approach to business service design and application architecture simply will not work, however. The drawbacks of cloud can heavily offset its benefits — most notably the unresolvable limitations of simple physics. Scientists have not yet been able to overcome the speed of light, and Einstein suggests they never will. Some other solution must be developed that does not depend on brute-force throughput over relatively slow and inconsistent WAN links.

That other solution involves more sophisticated architectures that sensibly locate compute and storage resources near one another and limit long-distance communications to more lightweight demands.<sup>6</sup> These principles have been understood for years and implemented in select situations, but cloud is forcing everyone to embrace them more. This is a good thing, but it will require extra work on the part of I&O professionals, application developers, and enterprise architects, among others. Focus more of this effort on newer systems of engagement and postpone legacy services and systems of record, maybe indefinitely. Rushing these older services into cloud models can prove downright destructive. I&O professionals should evolve to cloud quickly, but responsibly.

---

#### ENDNOTES

- <sup>1</sup> For more information on the Nirvanix collapse, please see the September 19, 2013, “[Quick Take: How Enterprises Should Approach Cloud Storage In The Wake Of Nirvanix Closure](#)” report.
- <sup>2</sup> For more information on cloud security, see the May 2, 2014, “[Brief: S&R Pros Remain Unprepared To Address Virtualization And Cloud Security Risks](#)” report.
- <sup>3</sup> For more information to help you decide where, how, and why to locate data center facilities, see the October 9, 2014, “[Strategically Locate Your Next Data Center](#)” report.
- <sup>4</sup> Equinix discussed Amazon Direct Connect at the 2013 Amazon AWS re:Invent 2013 conference. Source: “Network-Ready Your Hybrid IT Environment,” Equinix, November 15, 2013 (<http://www.youtube.com/watch?v=5cC0SmK5GkU>).

- <sup>5</sup> In C&S's deployment, the 22-minute to 8-second access reduction was found for accessing files (1.5 MB file) from California to remote workers in New York over a 10 Mbps WAN link. Source: Panzura (<http://go.panzura.com/rs/panzura/images/Panzura-CS-C%2BS.pdf>.)
- <sup>6</sup> Some of the more important principles of this new architectural approach are covered nicely in a recent Forrester report. See the July 25, 2014, "[Evolve Your Infrastructure Architecture For Systems Of Engagement](#)" report.

## About Forrester

A global research and advisory firm, Forrester inspires leaders, informs better decisions, and helps the world's top companies turn the complexity of change into business advantage. Our research-based insight and objective advice enable IT professionals to lead more successfully within IT and extend their impact beyond the traditional IT organization. Tailored to your individual role, our resources allow you to focus on important business issues — margin, speed, growth — first, technology second.

### FOR MORE INFORMATION

To find out how Forrester Research can help you be successful every day, please contact the office nearest you, or visit us at [www.forrester.com](http://www.forrester.com). For a complete list of worldwide locations, visit [www.forrester.com/about](http://www.forrester.com/about).

### CLIENT SUPPORT

For information on hard-copy or electronic reprints, please contact Client Support at +1 866.367.7378, +1 617.613.5730, or [clientsupport@forrester.com](mailto:clientsupport@forrester.com). We offer quantity discounts and special pricing for academic and nonprofit institutions.

---

## Forrester Focuses On CIOs

As a leader, you are responsible for managing today's competing demands on IT while setting strategy with business peers and transforming your organizations to drive business innovation. Forrester's subject-matter expertise and deep understanding of your role will help you create forward-thinking strategies; weigh opportunity against risk; justify decisions; and optimize your individual, team, and corporate performance.

« CAROL ITO, client persona representing CIOs

