

Product Brief

NitroAccelerator: SQL Server Replication Problems Solved

Date: April 2011 **Author:** Jon Oltsik, Senior Principal Analyst

Abstract: CIOs are quite busy these days, consolidating data centers to cut costs and building new applications to drive business processes. Both activities create a lot of data that often needs to be replicated across WAN links to feed real-time analytical applications, which can create application bottlenecks or demand an expensive bandwidth-hungry network. Enter [NitroSphere](#) and its NitroAccelerator for SQL Server WAN optimization product. When it comes to SQL Server replication, NitroAccelerator offers performance and cost benefits that make it the right tool for the job.


Overview

Over the past few years, CIOs faced a set of divergent objectives. On the one hand, they were asked to greatly reduce IT costs. To do so, many turned to data center consolidation and server virtualization as they moved remote systems to densely populated data centers. Amidst this “consolidate everything” mentality, however, CIOs were also tasked with building highly distributed applications in order to support global operations, enable local employees, and create a pool of business data for use in real-time analytics.

These opposing priorities are reflected in a recent ESG Research report.¹ When asked to identify the most important considerations for justifying IT investments in 2011, IT professionals pointed to economic factors like cost reduction and ROI as well as business drivers like supporting business processes and reduced time-to-market (see Figure 1).

Figure 1. Most Important Considerations for Justifying IT Investments

Which of the following considerations do you believe will be most important in justifying IT investments to your organization’s business management team over the next 12-18 months? (Percent of respondents, N=611, three responses accepted)



Source: Enterprise Strategy Group, 2011.

¹ See: ESG Research Report, [2011 IT Spending Intentions Survey](#), January 2011.

IT Priorities Lead To Data Replication Challenges ...

The combination of IT consolidation and distributed applications has led to massive database proliferation. Armies of databases residing in branch and regional offices were consolidated into larger regional data centers. At the same time, new database instances sprung up across the globe in local data centers, hosting facilities, or Internet POPs.

Database moves, adds, and changes were intended to support the IT initiatives described above, lowering costs, accelerating ROI, and to improve business processes. Unfortunately, these activities led to an unforeseen technical issue around database replication: with hundreds of databases now residing in large data centers, database replication between data centers consumed an enormous amount of WAN bandwidth. Simultaneously, global database replication requirements also presented challenges via their dependence on expensive and somewhat unreliable broadband connectivity. These issues may seem trivial, but WAN-based database replication problems can ultimately wreak havoc on data-centric business processes and the analytical applications they were designed to feed. The line of business manager will likely be enraged when a latency-sensitive control application can't receive data feeds from Thailand, Kazakhstan, and Ecuador needed to make real-time manufacturing and inventory decisions, placing business operations on [expensive] hold.

... And Most Technical Solutions Are Inadequate

Many IT professionals scoff at the database replication challenges described above, believing that there must be an easy fix. Simply tune the database, purchase more bandwidth from a global carrier, or inject some type of replication solution at either end of the network pipe—problem solved. Sadly, this isn't the case because:

- **Databases may not provide replication support.** Many databases are designed based on an assumption of cheap and bountiful bandwidth. For example, Microsoft SQL Server does not provide support for native Tabular Data Stream (TDS) compression. This means that CIOs must look elsewhere for help.
- **Replication needs alone may not justify bigger pipes.** Yes, database replication can eat up a lot of network bandwidth, but utilization can be extremely "bursty." The CIO may quickly veto a request for fatter network pipes when she realizes that aside from database replication, day-to-day networking activities only utilize 20% of total bandwidth capacity.
- **Broadband networks may be unavailable or unreliable in developing geographies.** Even if IT executives sign off on more bandwidth, it may not be available in remote regions of developing nations. And when it is available, broadband networks may have frequent service interruptions or travel through a multitude of network hops—clearly, a mismatch for real-time database replication needs.
- **WAN optimization solutions offer a shotgun approach.** Network-based WAN optimization can help accelerate database replication with TCP/IP compression, but it would be extremely costly to put an appliance alongside every database deployed across the globe. Furthermore, WAN optimization solutions are really designed to accelerate applications and files for remote users, not to replicate remote databases. Finally, WAN optimization solutions can be difficult to install, tune, and administer over time.

Without an elegant technical solution, most organizations are forced to come up with homegrown workarounds within their database or network infrastructures. These creative solutions can help speed up database replication in the short term, but are no match for network latency or additional data growth over the long haul.

Enter NitroSphere and NitroAccelerator

Like many other IT quandaries, database replication across the WAN is a case of having the right tool for the job. In this case, this means finding a solution that works in concert with the databases rather than the network that sits between them. Recently, ESG was introduced to a database replication solution called NitroAccelerator from NitroSphere of Austin, Texas, that fits this description.

NitroAccelerator is installed alongside SQL Server itself, intercepting packets at the network driver. Once installed, it optimizes and compresses the TDS protocol to provide LAN-like performance and latency over the WAN. Because NitroSphere "understands" the intricacies of TDS, it is able to ensure the integrity of the protocol is maintained while

optimizing its use of TCP/IP. This gives NitroAccelerator specific database intelligence in order to accelerate replication without sacrificing reliability or stability.

ESG recently spoke with a NitroSphere customer to gain a better understanding of database replication requirements and the NitroAccelerator solution. This particular firm had deployed a real-time analytics application based upon distributed SQL Server collectors located at Internet POPs across the globe. As data collection and replication requirements grew, the company looked for a SQL Server replication solution but couldn't find one that met its needs. The company was introduced to NitroAccelerator about two years ago and now uses the product throughout its network.

The best assessment of NitroAccelerator came from the customer itself:

- **On ease-of-use:** "All you have to do is install NitroAccelerator on each instance of SQL Server, where it inserts itself into the network stack on each server. Once this is done, it is very intuitive to use."
- **On performance:** "We ran tests before installing NitroAccelerator and then collected metrics after it was implemented. We saw a 75% to 80% TDS compression ratio and a 4x improvement in overall throughput."
- **On why it was needed:** "We are constantly seeing network connections go down, especially in rural areas of developing countries. This really accentuates the need to do database replication as quickly as possible."
- **On data center consolidation:** "We have been able to aggregate systems into large data centers in metropolitan areas that act as application hubs. This helped us save on data center fees and operating costs. We couldn't have done this without NitroSphere."
- **On cost savings:** "Our collocation facilities charge for bandwidth consumption. NitroAccelerator helped us contain these costs."
- **On business benefits:** "Some of our remote locations wouldn't be able to replicate without it. Others would be severely throttled, which could be just as disruptive to our applications."

The Bigger Truth

While data center consolidation and server virtualization continue to centralize IT, it is important to remember that other technologies will continue to drive the need for distributed network-based solutions. For example, the customer described above collects data around the clock, writes this data to distributed SQL Server databases, and then replicates the data to a central location to conduct real-time analytics. This type of use case has become commonplace for applications that track infectious diseases, measure global warming, optimize supply chains, or follow financial market activities.

Oftentimes, application developers and IT engineers build these applications with a myopic focus on data collection and analysis, forgetting to think about some of the more mundane technical requirements. Database replication fits into this category. Network bandwidth and WAN optimization can help here, but these are expensive, partial solutions at best.

NitroAccelerator doesn't provide sexy data visualization software or lightning fast network and server hardware for data analytics. It *does* provide a very strong engine for accelerating SQL Server across WANs, which serves as a key enabling technology for widely distributed "big data" applications. No one would accuse NitroSphere of being a name brand technology company, but its SQL Server acceleration prowess will likely make the company quite popular.