

# Keeping Your eCommerce App Online

## A Proven Architecture for Continuous Availability

### Introduction

eCommerce is big business, and it is quickly getting bigger. According to Forrester Research, eCommerce in the United States will top the \$300 billion mark in 2015 for the first time.<sup>1</sup> Forrester noted that eCommerce spending for the 2014 holiday season (November and December) totaled approximately \$90 billion, which equaled roughly one-third of the entire online retail bucket for the year. Cyber Monday in particular was a blockbuster, with the single day generating \$2 billion in sales.

Key drivers for this amazing growth are the increased penetration of mobile devices (including tablets), a greater share of wallet shifting to the web because of online shopping's convenience, increased spending from younger consumers, and the success of online merchants' marketing efforts.

With all this growth, shoppers are increasingly demanding an unparalleled buying experience: fast access via any device, available 24x7, with no page slow downs, content that's always accessible, and a frictionless checkout experience. The cost of anything less keeps going up. A 2013 outage at Amazon.com was estimated to cost the company \$66,240 per minute, according to Forbes<sup>2</sup>. Beyond revenue considerations, online retailers have to keep a close watch on competitor infiltration, the potential loss of future repeat business, negative brand equity, loss of productivity, and the ability of their supply chain to rapidly fulfill orders.

Online retailers must deliver an excellent experience at all times. The performance of the retailer's website is at the heart of this shopping experience. IT staff responsible for the online presence must not only keep the site up at all times but also swiftly scale for retail traffic surges. These traffic surges range from flash sales anytime in the year to peak season surges such as Black Friday or Cyber Monday. The lifeblood for the web application is the database, and unfortunately, in most IT shops, the database is the weakest link when it comes to system availability and agility. This critical infrastructure stack is difficult to keep always up, scale, and protect.

.....  
**eCommerce spending for the 2014 holiday season (November and December) totaled approximately \$90 billion, which equaled roughly one-third of the entire online retail bucket for the year.**  
.....

---

<sup>1</sup> Sucharita Mulpuru, Forrester – US eCommerce Forecast: 2013 To 2018, May 12, 2014

<sup>2</sup> Amazon.com Goes Down, Loses \$66,240 Per Minute, Kelly Clay, Aug 13, 2013

## Spending Money to Make Money

To support increased revenue and to drive innovation, eCommerce retailers are investing heavily for growth. Spending on eCommerce platform technology by larger U.S. companies will increase from \$1.204 billion in 2014 to \$2.090 billion by 2019, a compound annual growth rate of 12%. But manufacturers and wholesalers will increase their spending at a faster rate, as their share of eCommerce software spending will increase from 20% in 2013 to 30% in 2019, predicts the Forrester report “U.S. Commerce Platform Technology and Services Forecast.”<sup>3</sup>

Deploying an eCommerce platform is not cheap for the larger companies covered by this report. Forrester estimates the average project costs \$1 million, and 15% of implementation projects cost more than \$2 million. And that’s just the cost of the software: spending on implementation and maintenance services typically runs five times that of the software itself. Total U.S. spending on those services will also nearly double between 2014 and 2019, growing from \$4.200 billion to \$9.772 billion. Retailers investing in their eCommerce platform should ensure they’re also investing in improving the availability of the database infrastructure service.

## Challenges In Keeping Your eCommerce App Online

In addition to the cost and time of deploying an eCommerce platform, several other technical challenges plague the eCommerce infrastructure stack. These challenges include:

- Delivering 100% uptime – even through Black Friday and other traffic surges
- Enabling high performance – so customers don’t click over to your competitors
- Scaling on demand – to keep pace with sales growth and other traffic spikes
- Deploying security patches – so sites stay safe without downtime and lost revenue

### Delivering 100% Uptime

eCommerce operations is an unforgiving environment, where downtime is absolutely intolerable. To make eCommerce websites highly available and resilient to failures at the database tier, developers have had to build complex error-recovery logic to make sure database failures – and the errors resulting from them – could be appropriately dealt with such that consumers don’t see database errors, or the website doesn’t freeze or hang when a database failure or timeout occurs. The web application also needs to understand the database architecture and topology –

.....  
**To make eCommerce websites highly available and resilient to failures at the database tier, developers have had to build complex error-recovery logic to make sure database failures – and the errors resulting from them – could be appropriately dealt with such that consumers don’t see database errors.**  
.....

---

<sup>3</sup> Peter Sheldon, Michael Yamnitsky, Forrester – US Commerce Platform Technology And Services Forecast, 2014 To 2019, Feb 5, 2015

such as which server is primary and which is secondary, as well as discovering when it is safe to fail over between them. This type of logic becomes even more complicated to build within the eCommerce app when database architectures evolve to new replication types or require cross-location failover.

## Enabling High Performance and Scale On Demand

For a long time, database scalability and performance engineering has been a specialized art reserved only for the most demanding apps – where you could throw exotic hardware and software at the problem or for companies and teams that could dedicate the engineering and development resources required to make the applications more scalable and performant. This model is not only expensive and resource intensive – it's also very time consuming.

## Deploying Security Patches

The number of extreme-risk database and operating system exploits that require urgent patching has been on the rise over the last several years. Recent examples such as the SSL Heartbleed bug or the Shellshock exploit were extremely risky as they left almost every server in many large networks vulnerable and were exploited by hackers almost immediately at mass scale and used to spread malware, which led to many other attacks.

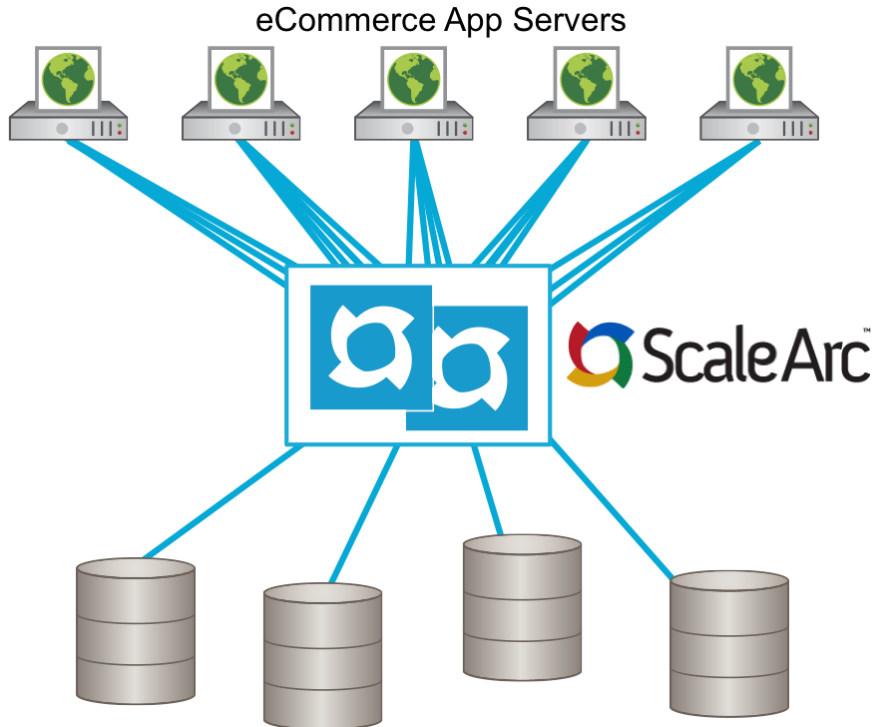
As a result of the increasing breadth and rate of vulnerabilities, security teams have mandated that patches to alleviate such exploits should be installed as soon as possible. Consequently, you likely need more frequent, almost ad-hoc maintenance activity that may require you to bring down database servers and, in turn, the eCommerce application. This type of maintenance activity is complex and tedious and requires resources outside of normal business hours. It may also require application re-deployment to add alert messages and maintenance window pages, it takes time to perform the actual maintenance operation, and then you must ensure the environment is healthy after the maintenance operation so you can bring up the app again.

.....  
**As a result of the increasing breadth and rate of vulnerabilities, security teams have mandated that patches to alleviate such exploits should be installed as soon as possible. Consequently, you likely need more frequent, almost ad-hoc, maintenance activity that may require you to bring down database servers and, in turn, the web application.**  
.....

## A Proven Architecture: Database Load Balancing Software

ScaleArc's database load balancing software deploys transparently between your eCommerce app and your SQL database infrastructure. ScaleArc leverages your database's replication and failover to enable continuous availability of your eCommerce app. It looks to the eCommerce app like a database and to the database like your eCommerce app. The ScaleArc software deploys on bare metal, VM, or in the cloud and offers a range of SQL management capabilities that dramatically improve the availability and performance of your eCommerce website, including:

- Failover that's transparent to the app – queuing transactions during failover
- Zero-downtime maintenance – patch on demand
- Scale out with no app changes
- ACID-compliant caching to boost response times
- Centralized point for access control and logging



**Figure 1: ScaleArc's database load balancing software deploys transparently between your eCommerce app and your SQL database. ScaleArc leverages your database's replication and failover to enable continuous availability of your eCommerce app.**

### Failover That's Transparent to the App

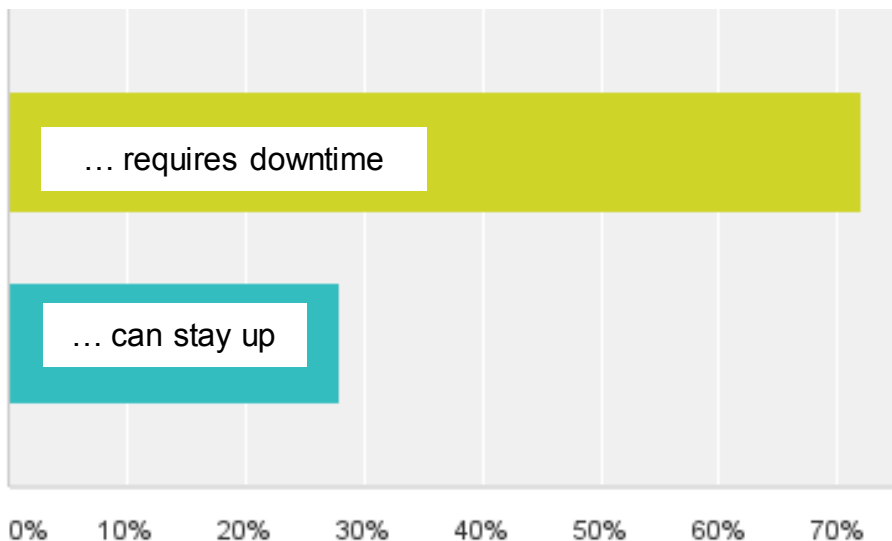
ScaleArc acts as a transparent abstraction layer, where the apps see ScaleArc as one universal database connectivity endpoint – one highly available connection string – that reduces complexity at the eCommerce app layer. ScaleArc then in turn knows the database architecture or topology and can detect and perform database failover in a seamless manner by employing its query queuing technology to shield the eCommerce app from unnecessary errors while database failure is occurring. If you're using failover technologies built into the database such as Microsoft's SQL Server AlwaysOn or Oracle's Active Data Guard, ScaleArc can interoperate with these technologies to prevent errors while the failover occurs. With ScaleArc, the recovery happens as quickly as the failover completes, without needing to wait for a DNS refresh or needing to establish new connections to the database.

If the database itself doesn't support automated failover – such as with SQL Server's transactional replication or MySQL's built-in replication or conventional Data Guard replication with Oracle – ScaleArc's auto failover APIs allow ScaleArc to orchestrate the process of database failover by executing the processes and commands to promote one of the secondary servers to take over as the new primary in a replication-aware manner. This approach allows for a significant reduction in Mean Time To Recovery (MTTR) for most eCommerce apps without the costs traditionally associated with more expensive database technologies such as Oracle RAC or SQL Server AlwaysOn.

## Zero-downtime Maintenance

Database maintenance is one of the most complex and disruptive activities for online retailers. Not only does this process require multiple teams to be involved in the operation to ensure it goes smoothly, but it also typically requires planned business downtime. In a recent survey ScaleArc conducted of 200 enterprises<sup>4</sup>, 93% of respondents at eCommerce companies reported that they needed maintenance windows at least once a month, and 42% needed them one or more times a week.

This process routinely requires redeploying your eCommerce application to add alert messages or maintenance window pages, performing the actual maintenance operation, ensuring the environment is healthy after maintenance operation, and bringing your eCommerce app back up. Most organizations are used to scheduling maintenance windows outside of business hours to reduce the visible impact of a maintenance window on customers or users – more than 85% of eCommerce companies in the ScaleArc survey said they scheduled non-critical updates after hours. It isn't uncommon to visit an eCommerce website at 11PM on a Saturday and find a maintenance alert up on the application webpage. However, in an always on world of anywhere, anytime, anyplace, as a matter of smart business, this practice is no longer acceptable to the buying experience.



**Figure 2. Data from a recent ScaleArc survey shows that during maintenance on critical apps, the app requires downtime over 70% of the time.**

Unfortunately, this type of maintenance activity is not only complex and tedious but is also becoming more frequent now that keeping servers up to date on the latest security patches is becoming a mandatory requirement in most organizations.

Online retailers have used database infrastructure such as Oracle RAC, SQL Server AlwaysOn, or MySQL Multi Master technologies to simplify the process of performing such maintenance – reducing the time to recovery but not necessarily eliminating the need for a maintenance window.

<sup>4</sup> The State of Application Uptime in Database Environments Survey, <http://www.scalearc.com/SurveyResults>, March 2015

Whatever level of database scale out and replication you have, ScaleArc works with your database infrastructure and lets you create a simpler, more economical, highly resilient architecture for your eCommerce environment. ScaleArc can leverage failover in the database or orchestrate the failover itself, providing the added benefit of eliminating maintenance windows so you can perform database patching or updates without any downtime.

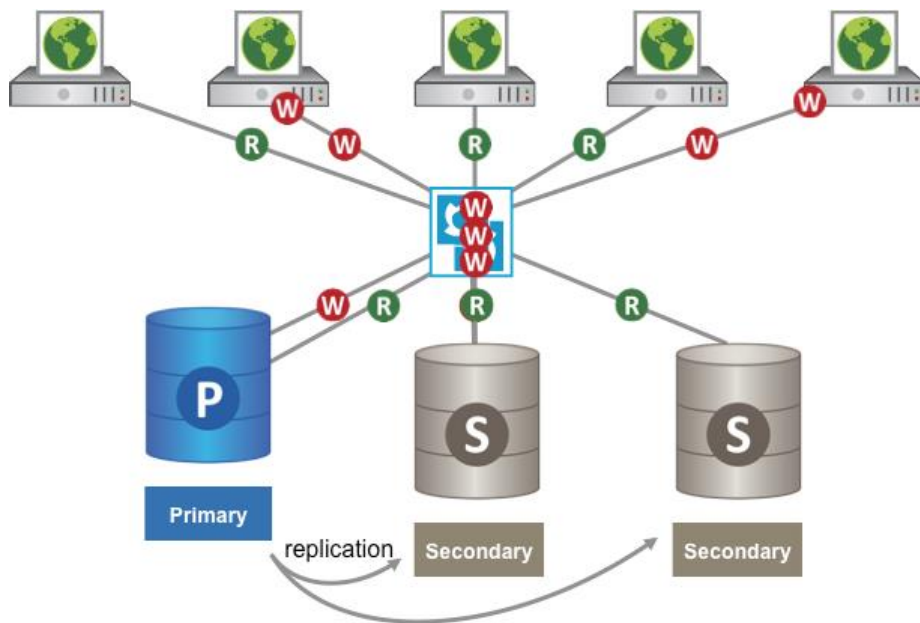
You can easily create a database cluster consisting of multiple VMs with ScaleArc and use ScaleArc as a control point for directing which database VMs are online and can receive traffic and which are down for maintenance. You can also failover between them without downtime thanks to ScaleArc's query queuing technology. This ScaleArc innovation ensures your eCommerce app database traffic can continue to flow even while you're performing maintenance on the databases, with no application redeployment or user-visible maintenance windows.

### Scale Out with No App Changes

ScaleArc simplifies database scalability and performance engineering, enabling any application and team to get a more scalable, highly performant database with just a few clicks. The ScaleArc features involved are:

1. Database load balancing and read/write split
2. Geo- and performance-aware load balancing

These features enable your eCommerce application to easily achieve significant improvements in database scalability and performance, without any application changes.



**Figure 3. By performing read/write split between the database infrastructure's primary and secondary servers, you maximize user experience on your site.**

### Database Load Balancing and Read-Write Split

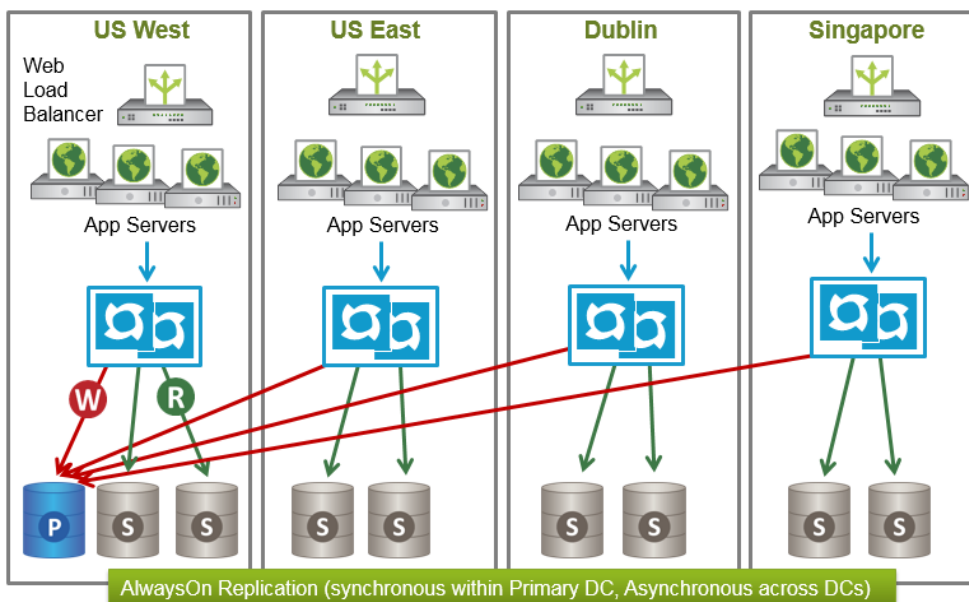
One of the main features of the ScaleArc software is its ability to seamlessly load balance the workload from your eCommerce application across a cluster of database servers. You may have had more than one database server in your

eCommerce environment already for DR or failover reasons, but very commonly, that database server does not serve database traffic unless the application is specifically modified to support reading from a secondary server.

ScaleArc makes it possible for your eCommerce application to use a whole cluster of database servers, without any application changes. Since ScaleArc is natively aware of the database protocol, it can identify what type of queries, stored procedures, or transactions are being run on a particular connection, and it load balances those queries in the appropriate manner. Imagine a database cluster with one primary or master server that can accept writes or transactions and two secondary or slave servers that can accept read-only queries. ScaleArc can automatically identify which queries are reads and which queries are writes or part of a transaction and will automatically send the writes or transactions to the primary server and load balance the reads across the available secondary servers.

### Geo- and Performance-aware Load Balancing

ScaleArc's load balancing algorithm is built to take various factors into account when distributing load amongst available servers in a database cluster. The three most prominent factors are replication lag, type of workload (whether read or write), and performance of the database servers itself.



**Figure 4. Microsoft's internal database environment for answers.microsoft.com. The setup spans four global geographies with ScaleArc deployed to augment SQL Server AlwaysOn.**

ScaleArc continuously evaluates the performance of each server in the database cluster by measuring the average Time To First Byte (TTFB) response time of all the queries sent to that server. ScaleArc then grades the servers in comparison to other servers with the same role – so all primary or master servers are compared against each other as part of the read/write pool and all secondary or slave servers are compared against each other as part of the read-only pool.



## ACID-compliant Caching to Boost Response Times

Complementing its Transparent SQL query caching technology, ScaleArc has introduced a new approach for cache invalidation to the industry – a method for automatically invalidating cache entries that enables true ACID-compliant caching. Data changes are tracked by extracting metadata from update or delete queries or from within SQL comments. With this invalidation method, ScaleArc can guarantee that stale data will not be served by its cache.

ScaleArc also provides the ability to handle more workload, reducing page download speeds and increasing website performance while protecting data. The following use cases are supported using the auto cache invalidation-based method:

### Shopping cart data

Shopping cart data is tracked for a user in an eCommerce application on every page load. Users browse through a lot of items before the final checkout, which puts significant load on the database. It is widely known that having slow load times impacts sales, and querying cart data from the database can slow web page performance. Up to 40% of web users abandon an eCommerce site if page load times exceed 3 seconds. Shopping cart data, despite being dynamic and unique for each user, can now be cached with auto cache invalidation since ScaleArc ensures data consistency and reflects the new items in ScaleArc's cache as soon as the cart is modified.

### Auction data

eCommerce sites where users can buy and sell items in an auction require the auction price and associated data such as current bidders to be fetched and updated frequently. The page load performance and the accuracy of auction data are extremely important for these sites. Database resources are heavily taxed toward the end of the auction when many more users are watching an item and starting a bidding war. Each item has a unique identifier that is used to query the data, making it a prime candidate for auto cache invalidation.

### User profile data

User profile data, that is personal data associated with a specific user, is typically the most accessed data. User profile databases need to have 100% uptime because they are validating users before any transactions can be performed. Users seldom modify their profiles, but when they do – such as update a password – your eCommerce app must immediately reflect the changed data. The user profile database lends itself very well to auto cache invalidation since each user has a unique ID within the database. Data tracking within auto cache invalidation keeps the user profile query cache up to date and significantly offloads the database from processing repetitive queries.

## Centralized Point for Access Control and Logging

ScaleArc is the only technology that can provide full, granular database access logging for all connections, reads, writes, and transactions without any additional performance overhead on the database stack. ScaleArc also processes these logs into very easy-to-understand usage patterns, which lets you analyze and compare data access patterns across a very wide time range very quickly and discover anomalies in a fraction of the time it takes to do so via conventional log-based



*"With ScaleArc, we achieved the best Black Friday in our company's history. ScaleArc ensured the database scalability we needed to support the launch of our Black Friday Survival Guide mobile application – without building any additional complexity into the application itself."*

*-- Craig Thayer, vice president of technology for Sazze*



analysis methods. Few other database logs would include such data anyway, since most databases store only write logs and not connection or read logs.

This ability to quickly compare access patterns is especially critical when a data leak or breach occurs, since the time it takes to discover the usage pattern of how the breach occurred is directly correlated with the time it takes to close the hole that led to the breach. The faster you close the breach, the less risk you face.

Though ScaleArc provides very detailed logs, the ability to store many months of data can be impacted if you're limited on storage capacity. ScaleArc's de-duplicated analytics conserve storage space by storing the analyzed, reduced-pattern data in easy-to-access, hourly patterns while still preserving many details about the database traffic, such as what type of query patterns were executed, how frequently, and by which users and what IP addresses. If storage is limited, the logs may be rotated, but this data is preserved.

ScaleArc's unique patterns comparison tool lets you quickly see what new query methods or patterns have been used in a time band when compared with another. You can easily find new query patterns that may correspond to SQL injection attacks or new, unique data-access patterns that have never been seen before. This tool lets you quickly find errant query behavior and block it using ScaleArc's database query firewall.

.....  
**ScaleArc is the only technology that can provide full, granular database access logging for all connections, reads, writes, and transactions without any additional performance overhead on the database stack.**  
.....

## Summary

eCommerce growth continues unabated driven by mobile penetration, a greater share of wallet shifting to online shopping, increased spending from younger consumers, and successful retailer marketing efforts. In the era of omni-channel commerce, retailers must successfully navigate the challenges of keeping their eCommerce website always on and high performing or risk losing consumers and the momentum that is driving revenue growth.

ScaleArc's database load balancing software deploys transparently between your eCommerce app and your SQL database infrastructure. Abstracting the database infrastructure away from your eCommerce application lets you failover transparently, patch on demand, scale out with no app changes, boost response times, and have a centralized point for access control and logging.

For more information about ScaleArc eCommerce Solution, visit [www.scalearc.com/eCommerce](http://www.scalearc.com/eCommerce).



*"At Virgin America, we see our website as a key touch point in the travel experience. ScaleArc's technology not only improves availability for the database servers, which underpin our website, but also offers impressive caching, acceleration and analytics capabilities to ultimately ensure the best web booking experience for our guests."*

-- Jake McClean, director, IT & Operations



2901 Tasman Drive, Suite 205  
Santa Clara, CA 95054  
Phone: 1-408-780-2040  
Fax: 1-408-427-3748  
[www.scalearc.com](http://www.scalearc.com)



ScaleArc is the leading provider of database load balancing software. The ScaleArc software inserts transparently between applications and databases, creating an agile data tier that provides continuous availability and increased performance for all apps. With ScaleArc, enterprises also gain instant database scalability and a new level of real-time visibility for their application environments, both on prem and in the cloud. Learn more about ScaleArc, our customers, and our partners at [www.ScaleArc.com](http://www.ScaleArc.com).

© 2015 ScaleArc. All Rights Reserved. ScaleArc and the ScaleArc logo are trademarks or registered trademarks of ScaleArc in the United States and other countries. All brand names, product names, or trademarks belong to their respective holders.

03/27/15