the 451 group

451 MARKET INSIGHT SERVICE | 2 May 2008 – IMPACT REPORT

ScaleOut scales up to address grid compute caching

Analyst: Matt Aslett
Sector: Enterprise Software

ScaleOut Software is one of a number of next-generation distributed caching vendors attempting to solve data bottlenecks with a caching middleware layer. The company has focused on Windows .NET environments to date but will introduce support for Linux and Solaris with version 4.0 in the next months. Version 4.0 will also feature a grid computing edition, which implements MapReduce-style parallel computing software framework functionality.

The 451 Take

ScaleOut lacks the brand awareness of its competitors but clearly has a technology that customers want. The fact that ScaleOut can lay claim to 160 paying customers indicates that there is significant interest in its .NET caching approach; this is especially impressive given its lack of marketing (the news section of its website hasn't been updated for more than a year). Support for mixed environments thanks to the delivery of Linux/Unix/Java functionality should further its adoption, while high-performance computing (HPC) and grid adoption provide good opportunity for expansion, especially in the lucrative financial services sector.

Context

ScaleOut Software was formed in 2003 by CEO and chief architect William Bain. Bain previously created a number of startups, including distributed Web load balancing software vendor **Valence Research**, which was acquired by **Microsoft** in 1998 and contributed the Network Load Balancing functionality to Windows Server.

ScaleOut does not disclose revenue figures but does break out an average sales price for the two industry segments it targets. For e-commerce server farms an average deal is $1,000 per server – typically with six to ten servers in a single deal. ScaleOut is also sold for use in HPC cluster environments. The company has only targeted industry sectors (as opposed to government and academia) to date for applications such as portfolio risk analysis, Monte Carlo simulations and algorithmic trading. In these environments, which can include anything from four servers to 10,000 servers, there is an average price of $2,500 per server.

ScaleOut has raised an undisclosed level of funding to date by Seattle-based **Technology Alliance's Alliance of Angels** group of individual investors, and is not currently looking for venture capital funding.

Products

ScaleOut's flagship product is the ScaleOut StateServer (SOSS). Like other distributed caching products, SOSS provides a middleware layer for fast-changing data, thereby offloading data processing and storage from the live database server. The end result is to

increase the performance of the database server while enabling data to be distributed across clusters of commodity servers, potentially decreasing infrastructure costs.

The product has primarily been employed in areas such as e-commerce, where it is used to increase the performance of online applications by providing a distributed cache for the data accessed between transactions. The company is now seeing increased adoption in HPC environments, where it reduces the need for message passing interface communication between the computing grid and back-end data servers. Additionally, ScaleOut also offers ScaleOut SessionServer, which is designed specifically to store ASP.NET session-state data – such as e-commerce shopping cart information.

ScaleOut GeoServer is an option for StateServer that extends the storage across geographically distributed server farms, while ScaleOut Remote Client enables remote access to the distributed cache. This means that StateServer can be deployed on a dedicated server farm accessed by any applications on the network.

With version 4.0, due for release in May, ScaleOut will also extend its reach with the delivery of a grid computing edition of StateServer. The company is looking to build on interest in HPC grid environments, in particular in the financial services industry.

Technology

ScaleOut's focus to date has been on the opportunity for providing distributed cache functionality for applications based on Microsoft's .NET. However, the company maintains that SOSS has always been architected to be portable across languages and operating systems.

With version 4.0, support for Linux, Unix and Java applications will reach general availability. Support is available for **Red Hat** Enterprise Linux and Fedora, as well as **Sun's** Solaris, while SOSS 4.0 also includes two new Java APIs to connect to Java applications that correspond to existing C# APIs for .NET. The Linux/Solaris support means that both Windows and Linux/Unix servers can be used to form a single distributed cache, and all servers can be managed through a Windows or Linux/PHP Web-based management console.

Meanwhile, the introduction of a grid computing edition is the result of an evolution of the SOSS technology since its introduction in 2005. From providing session-state storage in 2005 and application caching in 2006, the technology evolved to support data grid deployments for cache querying in 2007 and now compute grid deployments, including parallel methods similar to **Google's** MapReduce environment. This enables applications to more easily process large amounts of distributed data. The theory is that massive data sets can be processed across a distributed server farm by implementing two computations: map, which creates intermediary data sets; and reduce, which creates the smaller data sets, which are distributed across the server farm.

ScaleOut's Bain was involved in similar research projects at **Intel** during the 1980s and early 1990s – specifically, parallel domain decomposition. He describes ScaleOut's implementation as in-memory MapReduce, but is quick to add that the company is using its own proprietary merge algorithms and is not actually making use of Google's MapReduce algorithms.

Customers

ScaleOut boasts 160 paying customers, including **Reuters**, which uses StateServer to cache the headlines and first lines of new stories; the **Home Shopping Network**, which uses it to support its e-commerce shopping carts; and **CBS Radio**, which uses it to provide streaming media.

While the company's initial success has been in e-commerce applications, it is seeing more interest in HPC and grid environments, particularly in the financial services sector, which is dealing with huge growth in data volume.

Competition

Given ScaleOut's focus to date on Windows .NET environments, there is good reason why it claims to not see much of the early pacesetters in distributed caching in competitive situations. **Tangosol** (now **Oracle** Coherence), **GigaSpaces**, and **GemStone Systems** all have a background in Java applications and have only added support for .NET relatively recently. Although ScaleOut is now seeing more of Coherence and GigaSpaces, it also claims differentiation thanks to the fact that SOSS was designed from the start to be deployed on distributed systems, while rival products have scaled up from roots in single-server caching.

Additionally, both GigaSpaces and Oracle have also focused their attention upmarket on data grid deployments, so they are likely to compete more directly with ScaleOut in the future. ScaleOut claims not to have come up against GemStone despite its addition of .NET functionality. GigaSpaces recently told us that GemStone has dropped off its competitive radar, although that may have more to do with GemStone's recent focus on the ultra-low-latency needs of trading, and management of market data. Additionally, **IBM** introduced WebSphere Extended Deployment in April 2007 but has yet to make a big splash outside its own accounts.

In terms of direct competition, NCache from **Alachisoft** is a distributed in-memory object cache for .NET environments. ScaleOut admits that NCache's pure .NET environment has some features it lacks, but points out that the product lacks portability and support for heterogeneous environments.

At the other extreme, ScaleOut says it has never competed directly with **Terracotta**, which is an open source Java virtual machine clustering and cache provider. Despite the different technology backgrounds, ScaleOut is keeping an eye on Terracotta, however.

Strengths	Weaknesses
The fact that ScaleOut StateServer has been architected from the start for mixed environments and can support distributed caches on Linux, Unix and Windows servers should make it an attractive proposition for heterogeneous customers.	The focus on product development rather than marketing is laudable, but the company does need to pay more attention to its public image. It is a small company and still needs to invest in building out its headcount.
Opportunities	**Threats**
The adoption of HPC and grid computing in financial services is a significant opportunity, and ScaleOut also admits that it has yet to look at the government and academia sectors.	In Oracle and IBM, ScaleOut faces significant long-term competition, while GigaSpaces and GemStone also have a lead in terms of market positioning at the high end. There is a danger ScaleOut could find itself overshadowed.