

# A-TEAM GROUP

News, Research & Events for Financial IT Executives

[HOME](#)

[A-TEAM INSIGHT](#)

[RESEARCH](#)

[EVENTS](#)

[ADVERTISE](#)

[CONNECT](#)

[ABOUT](#)

- [About A-Team Insight](#)
- [Market Data](#)
- [Electronic Trading](#)
- [Low Latency](#)
- [Reference Data](#)
- [Risk-Technology.net](#)
- [A-Team Insight \(monthly PDF\)](#)
- [A-Team IQ \(quarterly magazine\)](#)
  
- [About Us](#)
- [Meet the Team](#)
- [Press Releases](#)
- [Contact Us](#)
  
- [About Research](#)
- [All Research](#)
- [Market Data Research](#)
- [Electronic Trading Research](#)
- [Low Latency Research](#)
- [Reference Data Research](#)
- [Risk-Technology.net Research](#)

<http://www.a-teamgroup.com/article/xtremedata-noetic-to-provide-the-edge/>

Article printed from A-Team Group (<http://www.a-teamgroup.com>)  
By [justin.magruder@noeticpartners.com](mailto:justin.magruder@noeticpartners.com) at 20:34:45 on 07 Mar 2011

## **XtremeData, Noetic Provide The Edge**

07 Mar 2011 [No Comments](#)

Data management specialist XtremeData and information architecture consultants Noetic Partners have launched The Edge, an integrated system for collecting, storing and performing analytics on massive volumes of market data.

What differentiates The Edge from existing analytics platforms, says Noetic's president and managing director Justin Magruder, is that it has been designed to store the market data universe in its entirety rather than a subset, which has traditionally only been possible, and which has been a big compromise for those running analytics.

The scalability of The Edge means that there is no need to decide what elements of data to keep. Magruder says the approach is to "keep everything ... there is no reason to throw out the bottom of the book."

The Edge comprises market data feed handler components, referred to as Listeners, running on multi-core, multi-socket AMD x86 servers, and XtremeData's distributed dbX database engine, which combines a massively parallel compute architecture with open source software and storage, providing scalability up to several petabytes of storage. With The Edge, real-time data feeds can be consolidated, loaded and fully indexed, and made ready for queries in 30 to 40 seconds.

Currently, says Magruder, Listener components have been created for Active Financial's ActivFeed, NYSE Technologies SuperFeed, and Thomson Reuters RDF. He expects further Listeners, including for direct data feeds and FX sources, to be built in due course. Overall, the architecture can handle up to 50 million messages per second, he adds.

As market data is loaded into dbX, additional fields are added, including 'noetictime' which normalises records from different markets for time synchronisation. Other fields include 'noeticsymbol' to uniquely identify an instrument that might trade in different venues.

In a typical user configuration, The Edge might comprise 100 terabytes of storage (from EMC or Hewlett-Packard), four Listener servers, each with four 12-core AMD CPUs and 64 gigabytes of memory, and a 10-node dbX database engine. Each dbX node - a commercially available blade from the likes of HP - itself comprises a 4-core AMD CPU, an Altera Stratix IV FPGA, 32 GB of memory, and 12TB of directly-connected storage, with nodes connected via 40 gigabit InfiniBand networking.

XtremeData's vice president of sales and marketing, Geno Valente, says a single dbX system can scale linearly up to 1024 data nodes, and 12.3 PB of storage. Beyond that, systems can still be connected together, although some care must be taken to maintain the linear scale characteristics.

Running on the dbX hardware is a Linux variant and a heavily modified Postgres database, which takes advantage of the distributed architecture for executing queries as well as leveraging FPGA hardware acceleration to run specific database access functions, providing an industry-standard SQL interface with very high performance.

Queries can be constructed to focus on individual instruments, or baskets, over different periods of time, performing calculations on the fly. As a demonstration example, querying symbol "C" - Citigroup - for all trades and quotes across all venues for an entire day, extracted 5.4 million records in around seven seconds.

Applications for The Edge include back testing of trading models, and market surveillance. Surveillance applications include searching for patterns of trading, and trade strategies, and their impact on the market - a particular focus since last May's 'flash crash'.

The Edge is available now from XtremeData, says Valente, with Noetic providing consulting for specific customisation. The system is currently being trialled at financial markets customers, with production implementations expected soon, he adds.

Categories: [0-Service](#), [A-Team Insight](#), [MDI-Pre-Trade Analytics](#), [Market Data Insight](#), [RRIT-Market Surveillance](#), [Risk & Regulation IT](#)

Article printed from A-Team Group (<http://www.a-teamgroup.com>)  
By [justin.magruder@noeticpartners.com](mailto:justin.magruder@noeticpartners.com) at 20:34:45 on 07 Mar 2011

---

© Copyright A-Team Group. All rights reserved.