

Clearing Postgres Misconceptions Among SQL Server Users

An EnterpriseDB
Position Paper

For DBAs, Database
Architects & IT Directors
June 2014



Table of Contents

Introduction	3
False Assertions Exposed	3
Open Source Truths Revealed	9
Conclusion	13
About EnterpriseDB	14

Introduction

PostgreSQL has advanced significantly in enterprise features and capabilities with recent releases and companies worldwide have accelerated their adoption of the database as a low-cost alternative to traditional databases. EnterpriseDB's Postgres Plus Advanced Server has also advanced its enterprise-class performance, security and manageability features. With this rising interest come more and more queries from Oracle, DB2 and SQL Server users exploring PostgreSQL (known simply as Postgres) for their database needs. Microsoft SQL Server users in particular have been very active in 2014, possibly because of recent changes in Microsoft pricing policies, or like so many others, have determined they are paying too much for their database.

The nature and consistency of queries from SQL Server users suggest a fundamental misunderstanding of Postgres features and capabilities. This paper seeks to inform SQL Server users about Postgres and correct misinformation that may have been targeted at this population. Below we will address specific claims made by the company and provide some additional insight to consider.

False Assertions Exposed

Open source software is a disruptive force, so it is understandable that existing proprietary vendors wish to seed doubt or concerns about its effectiveness. Given the rapidly expanding use of Postgres in particular, it is important to take a closer look at some of these assertions. After all, among just EnterpriseDB (EDB) customers are 50 of the Fortune 500 list of largest US companies and 98 of the Forbes Global 2000 list of the largest companies in the world – all using

Postgres.

If assertions were true that using Postgres would put operations at risk, then major corporations like ABN AMRO Bank, Grupo BBVA, Deutsche Börse AG, Ericsson, Fujitsu, KT Corp., Lockheed Martin, McKesson Corp., Nippon Telegraph & Telephone Corp. (NTT) and RSA Security LLC, as well as agencies throughout the U.S. government, including many within the Department of Defense and NSA, would not be using Postgres.

The growing number of large enterprises deploying Postgres for mission-critical applications is proof that Postgres is, in fact, a high-performance, secure and robust database with features and capabilities every bit as capable as the more expensive commercial alternatives. Not to mention more cost-effective.

The following is a series of common misconceptions about Postgres among SQL Server users, followed by the reality:

Misconception No. 1

Open source software ends up costing you more!

Feature comparisons between Postgres and SQL Server have suggested there are benefits to SQL Server in lowering costs compared to open source. Further, such arguments have posited that open source-based software is inferior because people who actually write the code don't have a commercial interest in the finished product.

It is correct to point out that there is no such thing as "no cost" when it comes to deploying any new software. There will always be costs for expert technical support, consulting, training and related tools and services regardless of whether the software is open source or proprietary. Further, many large enterprises work with companies like EDB to provide support, expert services and Postgres software that meets enterprise-grade demands for performance, security and manageability. Few organizations go it alone with just in-house experts.

But the fact remains that open source-based software is inherently less expensive than proprietary software. The open source development model does not have the large expense of development costs that

proprietary commercial software companies must recoup through higher prices. By having to fund just tools development and targeted enhancements, companies like EDB can begin pricing the database at price points 10 times lower than proprietary software providers.

Misconception No. 2

Open source code is more bug-ridden because the skills of the developers are unknown!

Proprietary software vendors would have you believe that anyone can introduce bug-ridden code into an open source project and that only a commercial enterprise can prevent this from happening by assessing a programmer's skill and adequacy.

Truth is, the Postgres code base is managed by multiple classes of contributors who serve in clearly defined roles. The core code base is protected by a team of long-term volunteers who are highly skilled and experienced community members called 'committers'. These individuals have been vetted and appointed by their peers, who are equally experienced. Only these most qualified people actually are permitted to check in code. They inspect, analyze and test code before it's accepted, and they reject code that does not meet the high standards of the PostgreSQL project including documentation and test cases. There are also redundancies to the process. Even top-level committers cannot check in code without a peer review.

The purposeful misrepresentation of open source development as a free-for-all is false and only serves the messenger with a vested interest in squashing competition. It's also quite disingenuous when the criticism comes from commercial vendors with reputations for releasing buggy and insecure products.

Misconception No. 3

It's impossible to know whom to contact about security patches in open source software!

Now that we've established there is structure and organization to the PostgreSQL community, it should come as no surprise that security issues are also managed in a disciplined manner, when they do arise.

Postgres, in fact, has a reputation for being the most secure open source database. It publicly reports and repairs security issues. The community does this primarily through the Common Vulnerabilities and Exposures organization. But the community also works cooperatively with 'packagers' of PostgreSQL, like EDB and other companies with ties to Postgres, to expedite patches to their respective user bases.

Anyone can access the 'Security' link on the PostgreSQL.org home page to report or view security issues. Try searching <http://cve.mitre.org/> for 'PostgreSQL' and 'SQL Server'. The reality is that open source projects with very large communities like Postgres are actually more capable of rapidly identifying and fixing potential security issues because they are unencumbered by corporate agendas and bureaucracies. Further, with thousands of regular contributors and hundreds of active developers, the open source PostgreSQL community has the manpower to respond rapidly.

Misconception No. 4

Day jobs get in the way of volunteer work so open source communities, without corporate program and project management, can't stay on track!

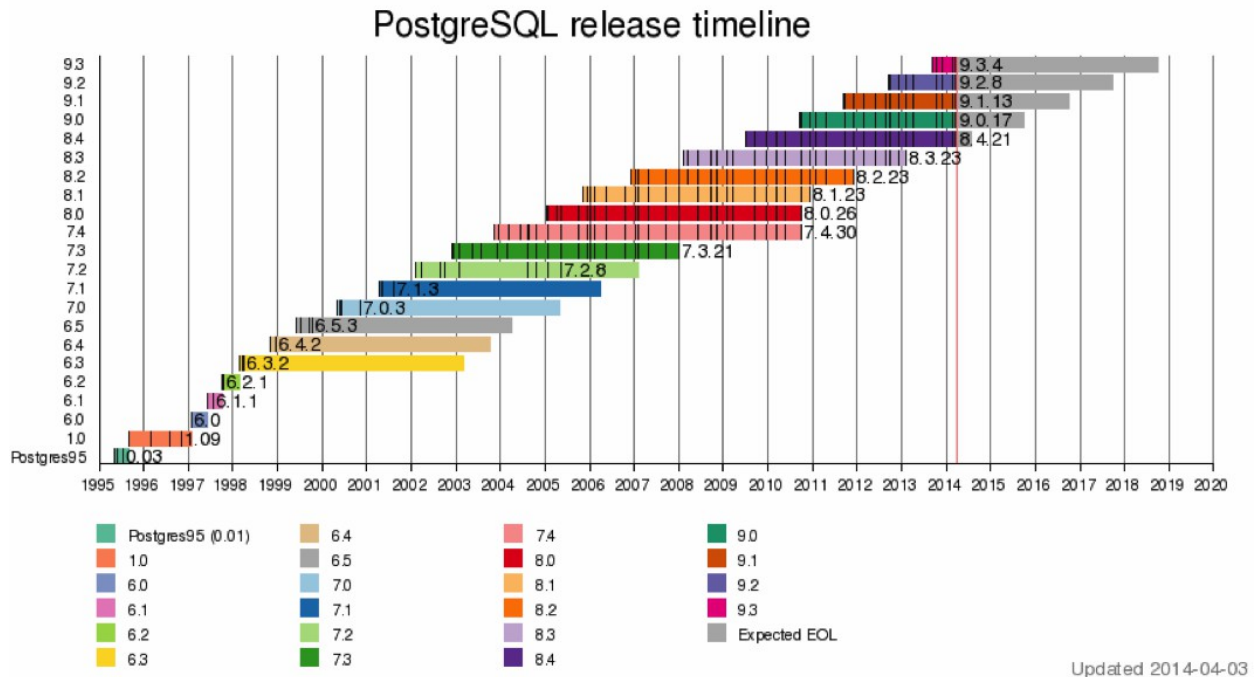
One of the great advantages of open source development is that the end user is often the developer too. This results in a greater dedication to purpose and quality among developers unparalleled within a traditional vendor. In addition, commercial entities like EDB bring enterprise-grade features into the project. EDB, in fact, employs more core community members, committers and other key contributors than any other commercial entity connected to Postgres. Because of this, the company subsidizes a great deal of Postgres development by supporting employee contributions at multiple levels.

As a community, Postgres actually functions quite efficiently thanks to the leadership structure, contributions by end users, and support from commercial entities like EDB. It has delivered increasingly high-quality, feature-full and market-heralded annual releases for the past four years. This includes multiple point releases as needed for the last four annual major releases. These activities have all run concurrently.

Another testament to the PostgreSQL community's project

Clearing Postgres Misconceptions Among SQL Server Users

management efficiency and effectiveness is its policy of supporting multiple versions of the software at once, which at the time of this writing is five but has been as high as six. The chart below shows a history of Postgres releases.



Now consider the case of the KNN capability to get a sharper sense of how well the PostgreSQL community functions compared to a corporate vendor. KNN is an acronym for k-nearest neighbors and is a search algorithm that has a lot of utility for a wide variety of applications. A team of scientists who needed KNN searches in Postgres developed the capability and made it available as an add-on plug-in feature in 2009 (See: <http://www.postgresql.org/message-id/4B0AC9E1.5050509@sigae.ru>). The code was further refined until it was in an acceptable form to integrate into the main code base, which happened in time for the 9.1 release in 2011.

Around the same time period, Microsoft was talking about releasing a similar feature. As of the time of this writing (May 2014) SQL Server had not yet deployed the feature. Instead, SQL Server users were

given the following workaround: write a query to find all the matches within a certain distance, sort results by distance, and if the number of matches is insufficient, expand the distance and try again. This feature allows you to ask for the N nearest matches (in physical distance or, for example, in name similarity) and that many matches will be returned from the index in distance order.

So much for corporate project and program management serving the needs of customers.

Misconception No. 5

With no corporate backing and all this community development, there is no specific lab environment to adequately test software in the open source world!

Another falsehood is that open source communities have no lab environments and no corporate backing, which ultimately leads to no hardware support. The Postgres community, in fact, has an extensive "build farm" of hardware and software platforms that are used to continually build the product and execute regression test suites. Each build farm component reports build and test results back to a server which the developers continually monitor to detect regressions or build failures on one or more of the wide variety of hardware platforms, operating systems and compilers in use in the build farm.

See http://buildfarm.postgresql.org/cgi-bin/show_status.pl.

Further, there is commercial backing behind these labs. Commercial interests maintain many of these machines, including EDB, which supports the Solaris SPARC, HP-UX, Itanium and Windows boxes.

Commercial backers that support the Postgres community in other ways include Red Hat, VMware, Rackspace, Google, HP, Salesforce and Skype. The full list of sponsors is available here: <http://www.postgresql.org/about/sponsors/>.

Misconception No. 6

Hardware vendors don't support open source databases so there are no compatibility assurances.

EnterpriseDB has, in fact, partnered with IBM and HP as well as Red Hat to ensure Postgres support on multiple hardware platforms. Most recently, EDB has worked with IBM to provide support for Postgres Plus Advanced Server on PowerLinux hardware. Finally, EnterpriseDB was awarded HP's 2013 HP AllianceOne Partner of the Year for Mission-Critical Computing as a result of a successful introduction of Postgres Plus Advanced Server to the German stock exchange Deutsche Börse.

In addition, the Postgres community as a group, as well as individual members such as those who work for EDB, work closely with a number of major hardware manufacturers to ensure that Postgres runs well on their platforms. This includes HP, Dell and Intel as well as other major players.

Open Source Truths Revealed

There are significant benefits to the PostgreSQL community functioning with no single sponsor that dominates or 'owns' the community. MySQL, for example, was originally dominated by a single sponsor, MySQL AB, which actually owned the code. Therefore, it could be bought and sold to the highest bidder, as demonstrated by its acquisition by Sun and then subsequently Oracle. Similarly, MongoDB is dominated by a single vendor. Postgres, on the other hand, is a truly independent open source project serving the broader needs of its users rather than a specific commercial entity.

Instead of a top-down corporate structure, the Postgres community is a meritocracy. Meritocracies are designed to allow the best ideas to win. In a meritocracy, everyone has equal access to information. Well thought out and compelling ideas will gain the most support and be the most successful. While many commercial interests strive to be merit driven, an open source project will quickly lose members and support when politics, power or money become the sole drivers for decisions. The longevity of the PostgreSQL project (18 years as it's currently structured) speaks volumes to the meritocratic personality and processes of the community.

Because it is open, the Postgres project functions with a great deal of transparency. One of the foundations of open source software is visibility into all aspects of the project - its source code, development processes, roadmaps, bugs, team members, etc. This kind of transparency gives users visibility into the quality, security, functionality, long-term viability and marketing honesty of Postgres.

Anyone can visit the community website, join the various mailing lists (visit: <http://www.postgresql.org/list/>) and just listen in on the whole process. Another component to Postgres' visibility worth mentioning is its excellent documentation available online for all versions. Are your database vendor's internal operations as visible as Postgres'? And if not, ask yourself: whom does that invisibility serve?

The community behind Postgres has amassed a number of achievements many corporations would find enviable, such as:

Cross-Platform Support

Postgres has extensive cross-platform support which allows for easy and more manageable integration of dissimilar IT infrastructures as is so often necessary when dealing with business partners, corporate mergers, acquisitions or experimenting with new technology. Multi-platform support provides greater flexibility in the IT infrastructure and greater leverage with software vendors since it is anathema to vendor lock-in. Companies often end up with a mixed bag of platforms—from acquisitions, partnership demands and evolution of IT. Because Postgres runs on nearly all platforms, including Linux, Windows, OS-X and others, enterprises benefit from seamless integration or transitions with their Postgres databases.

Strict Conformance to Standards

Postgres adheres to standards including the ANSI-SQL:2008 standard, rather than focusing on proprietary SQL variants that lock customers into vendor's products. There are benefit to this, including extensibility (which will be explored below) but also the portability of SQL skills back and forth among Postgres and other databases.

Platform Extensibility

Postgres is designed for extensibility as witnessed by: SQL/MED support (Management of External Data), extensible indexes, extensible

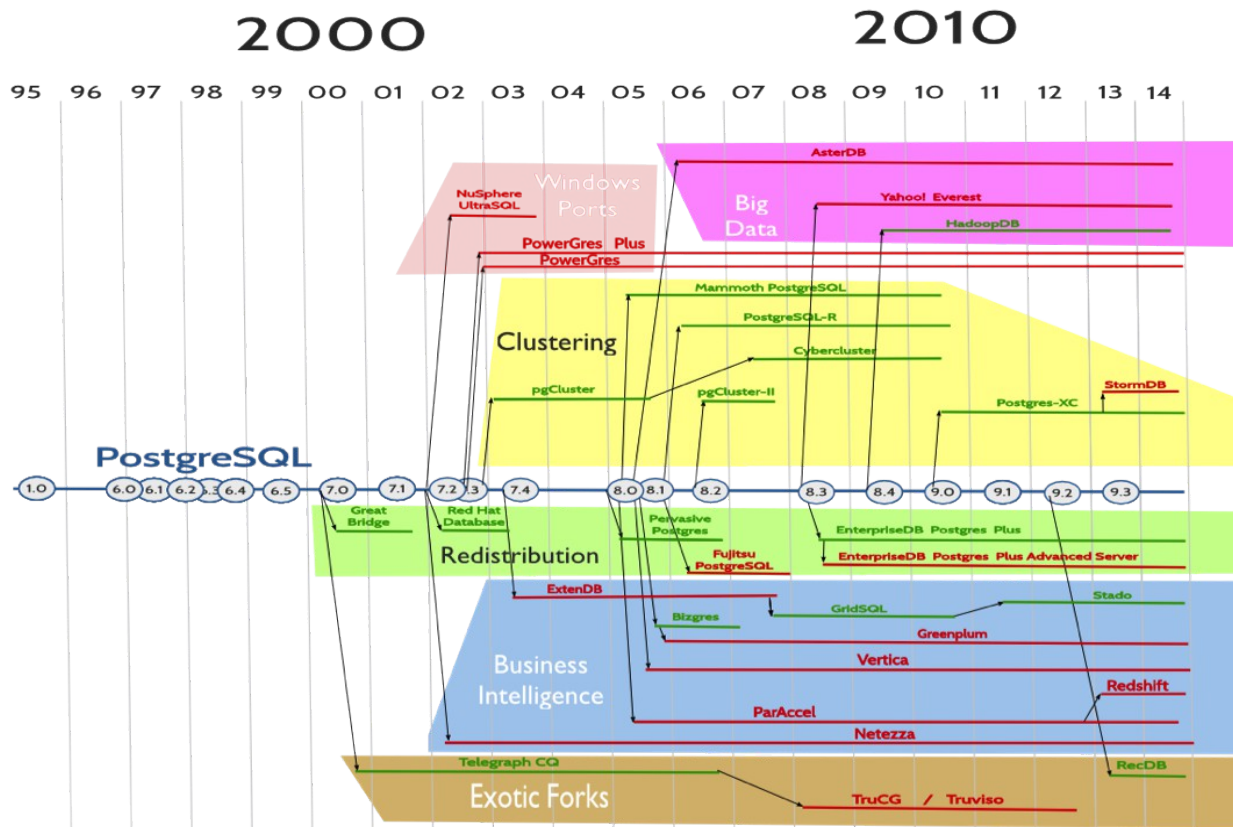
data types, operators, and functions, its support of more than 12 server-side programming languages, and its ability to easily integrate new data paradigms like key-value stores (HSTORE), and document data stores (e.g. JSON). Anyone can write an extension to the database. PostGIS is an excellent example. PostGIS is an open source plug-in for Postgres and is arguably the best geospatial database functionality in the market today among any relational database management system. It is also available at no cost. See: <http://postgis.net/>.

Another result of Postgres' extensibility is the wide variety of tools available in the community and from vendors like EDB. The community tools can be surveyed here: <http://www.postgresql.org/download/product-categories/>.

Expanding Ecosystem

Postgres comprises the foundation of the business intelligence industry as well as a number of embedded and black box database solutions. As there are many variations of SQL Server to serve customer needs, so too are there many variations of Postgres. These include VMware vFabric Postgres (from VMWare), Greenplum (from EMC) for big data analytics is the foundation of EMC's Big Data Division, Postgres-XC, a synchronous multi-master cluster designed for write-scalability, Postgres Plus Advanced Server (from EnterpriseDB) provides enterprise enhancements to Postgres plus database compatibility for Oracle, and HadoopDB from Yale university combines PostgreSQL with MapReduce technologies to address special large analytical workloads. The complete list of Postgres-derived servers is available here: http://wiki.postgresql.org/wiki/PostgreSQL_derived_databases.

These specialized versions of Postgres and their timelines are visualized below:



The Bottom Line of Open vs. Not

In revealing commissions of error and omissions of fact by a traditional database vendor, the hope is that end users can make up their own minds as to which database is better suited to their particular problems from a price-performance point of view. While not authoritative by any means, Wikipedia does provide a more unbiased public forum for the comparison of product features than those produced by a purveyor of one of the products. The following URL provides this comparison for many RDBMS's including PostgreSQL and SQL Server:

http://en.wikipedia.org/wiki/Comparison_of_relational_database_management_systems

When companies feel compelled to seek an alternative to any software they use, they must first ask a few fundamental questions. Of primary concern: Does the new software have the features I need? How much of my existing functionality can it replace? Will it save money in the long run?

If the alternative has the features you need, then comparing it to your existing software borders on distraction (not to minimize switching costs). If the alternative does not have all the features needed, then ask does it have enough features to address a broad enough class of applications that might still result in large cost savings and leverage with your traditional vendor?

EDB's Postgres Plus Advanced Server is a case in point for open source-based savings. With no large up-front licensing fees and a subscription-based purchase model (i.e. there are no CapEx expenditures) organizations purchasing Postgres Plus Advanced Server have seen savings as high as 80-90 percent of their previous database budgets.

Conclusion

Red Hat with Linux and other commercial entities supporting open source software have proven that the open source model can and does provide excellent performance and lower costs to organizations often with better feature sets than traditional vendors. Now EnterpriseDB with Postgres is disrupting the economics of database tier.

Many misconceptions and misinformation are being spread about Postgres particularly as more enterprises deploy the database. Yet the truth remains that the open source database Postgres performs on par with traditional databases for large swaths of applications across the enterprise at a price / performance factor that translates to impossible-to-ignore cost savings.

There is a place for Postgres in most if not all IT shops. The task then for DBAs, Developers, Database Architects, and CIOs is make sure that the right database is matched to the right application and in particular, at the best price / performance available.

EnterpriseDB's evaluation and migration assessment methodologies and experienced professionals will help you plan your migration. Contact us at +1-877-377-4352 or +1-781-357-3390, or send an email to sales@enterprisedb.com to get started on your migration.

About EnterpriseDB

EnterpriseDB is the leading worldwide provider of Postgres software and services that enable enterprises to reduce their reliance on costly proprietary solutions and slash their database spend by 80 percent or more. With powerful performance and security enhancements for PostgreSQL, sophisticated management tools for global deployments and Database compatibility for Oracle, EnterpriseDB software supports both mission and non-mission critical enterprise applications. More than 2,500 enterprises, governments and other organizations worldwide use EnterpriseDB software, support, training and professional services to integrate open source software into their existing data infrastructures. Based in Bedford, MA, EnterpriseDB is backed by top-tier venture capitalists and strategic investors like Red Hat and IBM. For more information, please visit <http://www.enterprisedb.com/>.

20140611

