

Data Warehouse Appliances: The Next Wave of IT Delivery

RISE OF THE MACHINES: APPLIANCES ARE THE NEXT WAVE OF IT DELIVERY

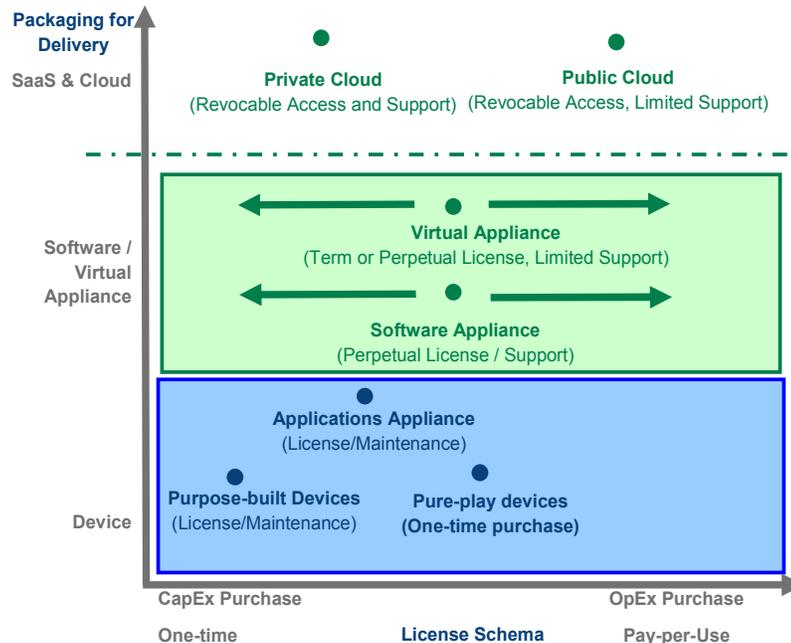
Large and small businesses must address the challenge of meeting ever-expanding workload requirements and rapidly delivering ROI while reducing execution risk on tightening budgets. Appliances – or purpose-built devices that pre-integrate hardware and software to address specific workloads – are rapidly becoming a preferred purchase option.

IT buyers face a growing field of choice as innovation increases the number of IT delivery options available to address specific workloads and as business evolutions increase a purchaser’s contracting options. Add in an increasing number of systems vendors such as IBM, EMC, Cisco, HP and Oracle who can create and sell broad portfolios of IT solutions as services, point products, bundles and appliances and the field of choice can appear to be overwhelming for IT purchasers.

Problem: IT buyers have to meet workload and service-level requirements while delivering ROI and reducing risk.

Complications: There more delivery mechanism vendors can use to package the functionality buyers require; and there are more vendors entering the appliance space, making buying decisions unclear.

The Expanding Field of IT Packaging and Contracting



SOURCE: TBR

Appliances are rapidly becoming a preferred purchase option for large and small businesses seeking to meet expanding workloads and deliver ROI in the face of tightening budgets. TBR is reporting the results of a 1,300 end-customer survey regarding their adoption of appliances.

In order to clear up confusion and assist buyers who are looking for a solution to their data warehousing requirements, TBR is reporting the results of a survey of 1,300 end-customers about their adoption of appliances. This white paper delivers peer-based insights to end-customers who are considering an investment to address their data warehousing requirements. The report:

- Defines appliances – or purpose-built devices – as a rapidly emerging opportunity for some buyers.
- Examines the nature and limitation of appliances over traditional custom integrations.
- Reports on key end-customer factors that are “must-haves” and “delighters” relevant for data warehouse appliances.
- Compares the key benefits of EMC’s Data Computing Appliance (DCA) offering against end-customer requirements.

APPLIANCES ARE NOT BUNDLES

TBR conducted a survey of over 1,300 corporate IT purchasers to determine the drivers of appliance adoption.

Over 50% of respondents indicated that they’d purchased at least one appliance in the past 12 months.

TBR defines a purpose-built device or appliance as a packaged and integrated platform of hardware and software, designed to deliver pre-defined functionality through a specific interface with no access to the underlying and supporting software.

Appliances are not bundles of hardware and software that require integration, testing and internal configuration.

APPLIANCE FORM FACTORS VARY – COMPLICATING THE PICTURE FOR CUSTOMERS AND VENDORS

Appliances vary widely in their form factors. TBR examined a wide array of different form factors, including:

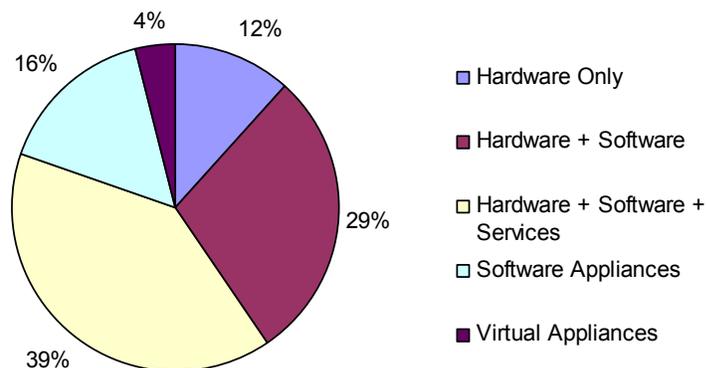
- **Hardware Appliance:** An optimized compute configuration
- **Hardware + Software Appliance:** Optimized compute, storage, network configuration and a management stack
- **Hardware + Software + Services Appliance:** Optimized compute, storage, network, management, OS environment, application stack, and set services
- **Software Appliance:** Optimized virtual software stack that runs either on commodity hardware or in the cloud
- **Virtual Appliance:** A virtualized software application with bundled hypervisor that runs either on commodity hardware or in the cloud

While combinations of hardware and software are the most commonly purchased form factor, virtual appliances are beginning to be reported among appliance purchasers. The increasing abstraction of hardware into pools of commodity compute and storage allows the software layer to be decoupled from the hardware. This trend benefits the purchaser and the vendors, as the virtual appliance can be deployed even more quickly and the revision or upgrade cycle can increase in speed. Potential buyers of appliances should look to vendors who can provide this layer of abstraction and efficiency as an option in their appliance delivery strategy.

Appliance Form Factor Purchased



What form factor did you purchase for the workload?



SOURCE: TBR.

Combinations of hardware and software remain the most commonly purchased form factor, but virtual appliances are increasingly being reported among appliance purchasers.

TBR LOOKED AT THE APPLIANCE MARKET THROUGH THE LENS OF WORKLOADS

One of the key differentiators of the study is that TBR examined the appliance market by looking at the workloads addresses by the appliance.

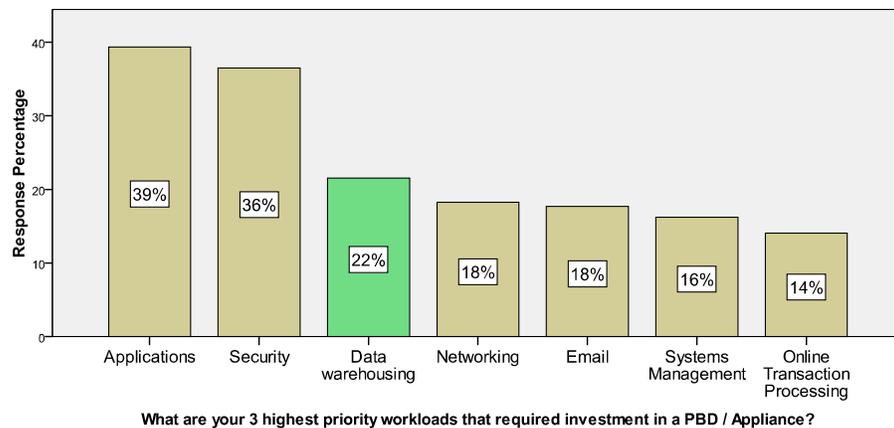
In the TBR survey of IT Decision-makers who recently purchased an appliance, data warehousing was the third-most frequently mentioned workload reported among recent appliance purchasers, with over one-in-five appliances purchased. Applications and security topped the workloads.

IT Workloads Tested:

- Applications
- Application Development & Testing
- Analytics & Data Mining
- Business Processing
- Data Warehousing
- Infrastructure Database
- Online Transaction Processing
- Security
- Systems Management
- Technical Computing

Respondents indicated that the increasing volume of data, the need to consolidate multiple sources of data and the requirement for extracting business insight from that data as the drivers for their purchase of a data warehouse appliance.

Workloads Addressed Through an Appliance Purchase



SOURCE: TBR

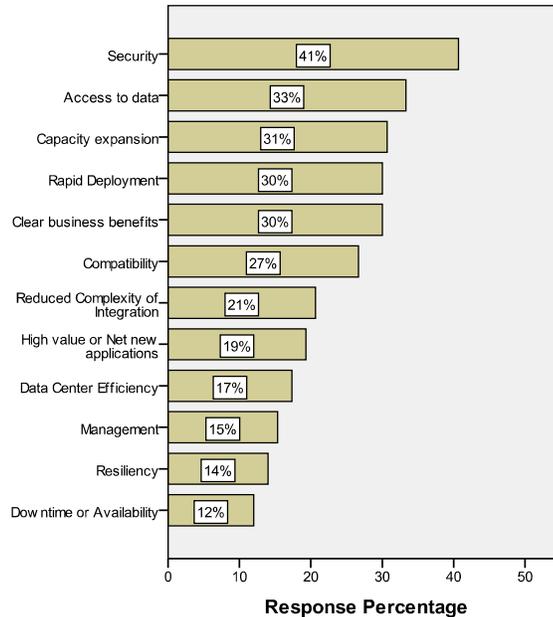
Survey respondents indicated that the main drivers for their purchase of a data warehouse appliance include the increasing volume of data, the need to consolidate multiple data sources and the need to extract business insights.

The data warehouse appliance market is one of the more mature segments of the appliance market space. Vendors such as Netezza pioneered the concept of an appliance dedicated to data warehousing, and other vendors rapidly followed. Today, nearly every large IT vendor sells data warehousing appliance space includes IBM, HP, EMC (Greenplum), SAP (Sybase), Oracle (Sun), Teradata, and Microsoft (DATAlegro):

Purchase Drivers for Data Warehouse Appliances

Security was the most frequently cited reason for the purchase of a data warehousing appliance.

Why did you purchase an Appliance for the Data Warehousing workload?



SOURCE: TBR

Purchasers of data warehouse appliances were asked why they purchased an appliance rather than purchasing component parts that require integration. The most frequently reported reasons include:

- Security of the data was mentioned most frequently as a key reason as pre-integrated data warehouse devices enable standardized security solutions.
- Rapid and easy access to the data and scalable capacity expansion are the next purchase drivers, as data warehouse appliances are optimized for their workload and offer incremental capacity expansion.
- Rapid deployment is the fourth most frequently mentioned key purchase driver, as data warehouse appliances ship as an integrated unit and do not require on-site build-test time.

TBR EXAMINED DATA WAREHOUSING APPLIANCES TO DETERMINE “DELIGHTER” AND CRITICAL “MUST HAVE” BENEFITS

Customers do not view all benefits through the same lens

Business benefits such as less training and rapid deployment are more likely to drive satisfaction, where technical benefits such as compatibility, functionality and security are key drivers of dissatisfaction if not implemented by the vendor.

Every workload-targeted appliance has its own mix of benefits that are “must haves” and that “delight” customers

Data warehouse devices carry a unique mix of perceptions across the benefits of appliances. There are two main groupings:

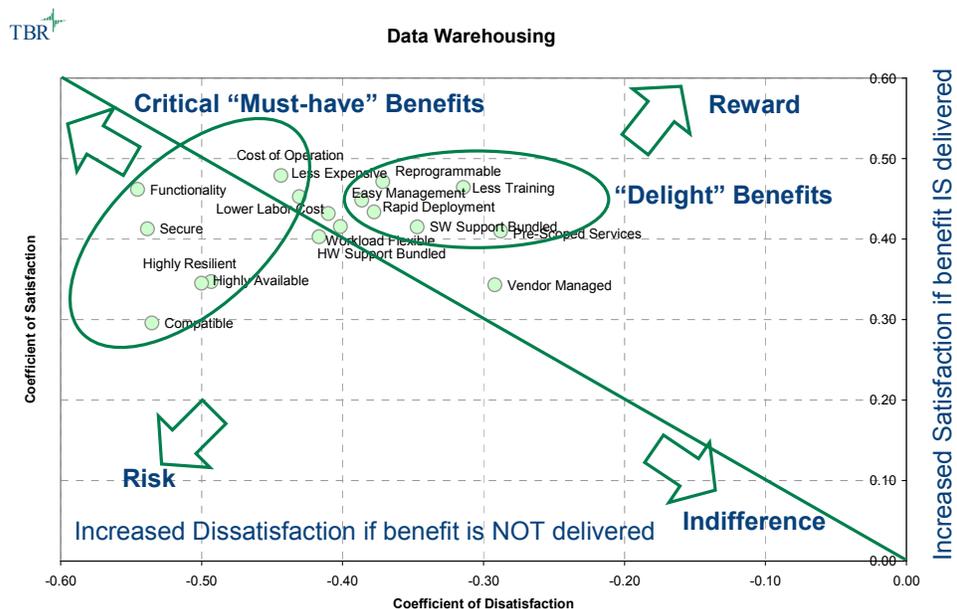
1. Critical “must-have” benefits that must be present or the purchaser will experience very strong levels of dissatisfaction; and
2. “Delighter” benefits that will deliver high levels of satisfaction if the vendor can deliver.

In terms of delivering a data warehouse appliance, customer will experience very high levels of satisfaction if the “delight” benefits are delivered by the appliance. In other words, these are soft benefits that arise from the pre-packaging and pre-integration of the appliance form factor.

Customers who are considering a data warehouse appliance should evaluate their options based on these two sets of benefits, differentiating between the functional, performance and cost-saving benefits in the “must-have” category from the secondary benefits associated with the “delight” category of benefits.

The benefits of data warehouse devices are grouped into either “must-have” or “delighter”

Map of Benefits Balancing Satisfaction and Dissatisfaction



SOURCE: TBR APPLIANCE ADOPTION SURVEY JAN 2010, N=696.

Key “Must-have” benefits include:

1. Specific Functionality (what the PBD / Appliance does)
2. Security/Data Protection
3. Cost of Operation (cost to operate is lower than custom solution)
4. Resiliency (PBD / Appliance functions if a component fails)
5. High Availability / Service level
6. Compatible (works with on-premise solutions)
7. Pricing (cost to purchase is less than total of components)

In terms of delivering a data warehouse appliance, customer demand the “must-have” benefits and vendors risk high levels of dissatisfaction with the appliance if these benefits are not delivered. In other words, these benefits are the most important benefits that the appliance form factor delivers.

Key “Delight” benefits include:

1. Reprogrammable (to add new functions)
2. Ease of Management (PBD / Appliance has a single management console)
3. Ease of use (offering requires less training than custom solutions)
4. Speed of deployment
5. Software support (annual support included in purchase)
6. Packaged Professional Services available to maximize ROI on device.

DATA WAREHOUSE APPLIANCE PURCHASERS MUST ENSURE THAT VENDORS DELIVER THE RIGHT MIX OF BENEFITS AS WELL AS THE RIGHT BALANCE OF PACKAGING AND CONTRACTS

The spectrum of options available for the delivery of IT benefits is expanding rapidly. From on-premise installation of IT purchased on capital budgets, customers now have a broad array of delivery options ranging from the traditional on-premise model through managed services, cloud services, hosting to a fully outsourced business process. At the same time, IT vendors are expanding the ways that customers can purchase IT, adding new options such as utility pricing, pay-per-use, leasing, term payments to the traditional capital purchase.

Customers are faced with having to balance the trade-offs between control over the IT assets that they received with the on-premise model with the cost, expertise, and scale benefits that can be achieved by shifting part of the burden to a provider.

The spectrum of options available for the delivery of IT benefits is expanding rapidly.

Appliances specifically balance the control received through an on-premise deployment with the cost and ease of use benefits created through the pre-packaging and pre-integration of the components. For example, vendors can deliver virtualized software appliances that are optimized to run on pre-integrated hardware platforms, selling the solution as a capital or lease option for the hardware and a subscription for the software.

Through TBR's research of enterprise purchasers of data warehouse appliances, we now understand that customers should look to buy a data warehouse appliance instead of a bespoke data warehouse solution under the following conditions:

When to purchase a data warehouse appliance:

1. Reduction in costs is more important than highly optimized "tuned" performance
2. Seeking reduced administration expenses through pre-integration and the single console
3. Requirements call for built-in high availability through hardware + software design or through virtualization and abstraction
4. When rapid and lower-cost scalability through modular, pre-designed and pre-integrated architectures is a requirement

Using these key "Must-have" and "Delighter" criteria listed above, customers can evaluate the offerings of the competing vendors to determine which data warehouse appliance best meets their needs and that will drive "delight" instead of increased expense, management effort and needless customization.

EMC'S "DATA COMPUTING" VISION PROVIDES FLEXIBILITY FOR EVOLVING CUSTOMER REQUIREMENTS

As customers focus on business as well as IT benefits of their data warehouse appliance, vendors are beginning to revise and update the packaging and benefits of their offerings. Following the acquisition of Greenplum, EMC launched a new Data Computing division that integrates EMC's best of breed backup and recovery solutions, Greenplum's shared-nothing, MPP analytical database technology and VMware's virtualization platform into a single data warehouse platform. EMC is attacking the data warehouse market by leveraging its core assets as well as by transforming data "warehousing" into "data computing." EMC defines data computing as a new data warehouse paradigm which moves processing dramatically closer to the data and analysis closer to the people who need insight. Data computing has the potential to be the next transformative step in data warehousing.

While data warehouse offerings have traditionally been seen as a powerful lens for the analysis of in-house structured data, EMC's "data computing"

TBR believes customers should consider a data warehouse appliance when specific conditions are met.

concept redefines the data warehouse to include external data sources in the analysis and integration with 3rd party business intelligence and analytical tools. The result is a device that targets the management and integration of data and information and the connection of the data to appropriate analytical tools. Instead of a pre-configured data repository, EMC offers a more dynamic device. Loading and query times – key buyer performance criteria – are simplified to enable easier end-user access to data and to improve speed and performance in the face of increasing data volumes.

EMC reports a strong value proposition for the Greenplum Data Computing Appliance (DCA) based on increased flexibility and global control of data management. Based on TBR research, security, resiliency, compatibility and high availability are key dimensions EMC can deliver that end-customers rate as “Must-haves” in their data warehouse device. EMC product dimensions that are “Delighters” include the two highest rated criteria of reprogrammability and ease of management (especially for internal and external sources of data).

Comparison of Key End-Customer Criteria with Greenplum Data Computing Appliance Value Proposition

EMC is attacking the data warehouse market by leveraging its core assets and by transforming data “warehousing” into “data computing”.

| Top Criteria | End-Customer Priority | Reported Greenplum Data Computing Appliance Value Proposition |
|--------------|---|---|
| “Must Have” | <ol style="list-style-type: none"> 1. Specific Functionality (what the PBD / Appliance does) 2. Security/Data Protection 3. Cost of Operation (cost to operate is lower than custom solution) 4. Resiliency (PBD / Appliance functions if a component fails) 5. High Availability / Service level 6. Compatible (works with on-premise solutions) | <ol style="list-style-type: none"> 1. Functionality – Offering spans corporate and external data sources 2. Includes RSA offerings 3. Unclear – EMC has not reported pricing 4. VMware virtualization offers strong resiliency 5. VMware virtualization supports HA – hardware HA is unclear 6. Compatible – works with third-party business intelligence offerings |
| “Delighters” | <ol style="list-style-type: none"> 1. Reprogrammable 2. Ease of Management 3. Ease of use | <ol style="list-style-type: none"> 1. SpringSource tools enable the addition of new functionality 2. The Greenplum Data Computing Appliance enables the integration of internal and external data in a single device 3. Once built and integrated, a DCA device that integrates across multiple third-party analytical tools should be easier to use than multiple data marts. |

SOURCE: TBR AND EMC



The emergence of DCA + business analytics systems represents a significant milestone in the evolution of computing.

TBR believes EMC is positioning the device to deliver net benefits to both business and IT users – a key requirement for success in the appliance market. Additionally, the device plays to EMC’s strengths in information management, virtualization and security. Finally, the device is agnostic at the application and analytics layer – providing flexibility to end-users who want to choose among vendors rather than standardize on a single vendor. TBR believes buyers considering a data warehouse appliance – especially those who can leverage existing EMC and VMware investments – should examine EMC’s DCA offering.

Technology Business

Research, Inc., is recognized as a leading independent technology market research and consulting firm specializing in the analyses of hardware, software, networking equipment, wireless, portal and professional services.

Serving a domestic and international clientele, TBR provides timely and accurate market research and business intelligence in a format that is uniquely tailored to clients’ needs.

TBR analysts are available to further address client specific issues or information needs on an inquiry or proprietary consulting basis.

For more information please visit www.tbri.com