

Technology Brief

## DH2i Has a Disruptive Approach to Protecting SQL Databases that IT Should Check Out

**Date:** March 2014 **Authors:** Jason Buffington, Senior Analyst; and Monya Keane, Research Analyst

**Abstract:** Most data protection and availability technologies are based at the server or storage level rather than higher, at the application level. A company called DH2i is taking a fresh approach to assured availability, one that recognizes our complete dependence on data and the pervasiveness of Microsoft SQL Server as an enabling platform. It is an availability and resiliency solution that SQL administrators should be looking into closely.

### Overview

IT vendors and IT pundits tend to toss around the word “disruptive” a lot—sometimes unjustifiably. But on occasion, a technology emerges that does work in a smarter way, leading to enough new IT-delivery opportunities that it becomes a disruptive force. Data protection is a segment of the IT industry boasting interesting candidates for a disruptive label, in part because of the variety of activities and efforts that modern data protection must cover.

Today, data protection is a “superset” term, with a core composed of backup, plus snapshotting, plus replication. The complete spectrum encompasses even more, including archiving and availability (see Figure 1).

Figure 1. The Spectrum of Data Protection



Source: Enterprise Strategy Group, 2014.

The spectrum is a reminder that data protection on its own has no value; the only reason data protection activities happen at all is to ensure subsequent accessibility. Such scenarios commonly relate to ensuring access to data after a crisis or outage. But *everything* in data protection boils down to “assured access.” For example, archiving provides assured access to older or business-critical data. Availability provides assured access without interruption—including continuous access to the data (at a minimum) *and* service availability: Having a VM running properly won’t matter if the front-line application can’t get to its underlying data.

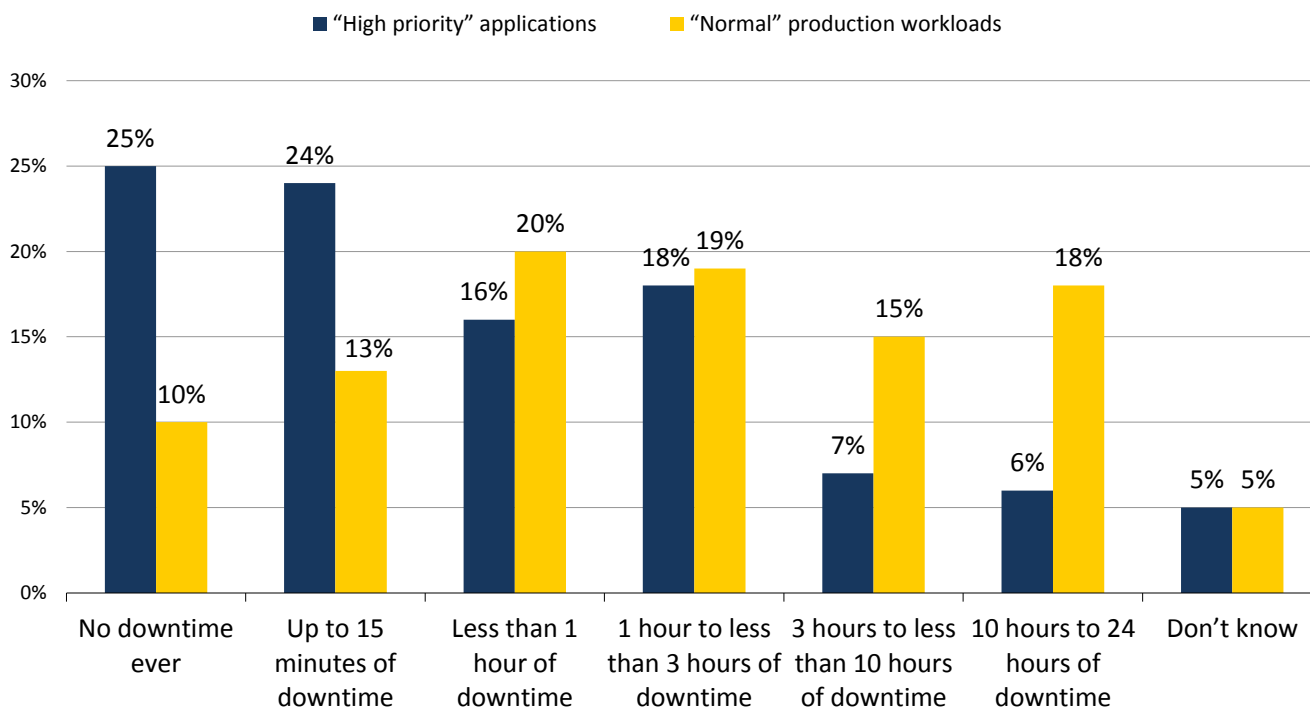
In general, data protection efforts today center on putting necessary *tactical* elements into place to facilitate *strategic* business capabilities—capabilities that are not limited only to disaster recovery and business continuity.

## Downtime Intolerance

End-users can barely stomach application downtime anymore. They demand assured access. More than a third (37%) of respondents to an ESG survey reported that downtime for their applications protected in the cloud cannot exceed an hour without causing an adverse business impact.<sup>1</sup> Even so-called normal, non-mission-critical apps have a very low downtime tolerance profile these days (see Figure 2), which actually represents a shift: Tolerance attitudes for tier-2 and tier-3 application downtime used to be far more permissive and lenient.<sup>2</sup>

Figure 2. Downtime Tolerance for Tier-1 and Normal Applications

**What are your organization’s RTOs for “high priority” applications and “normal” production workloads? (Percent of respondents, N=325)**



Source: Enterprise Strategy Group, 2014.

Why are SLAs for all business apps so tight now? Because we have become absolutely dependent on IT applications to do our jobs.

Oddly enough, most availability architectures are designed with barely any recognition of the actual applications involved, instead focusing on the availability of the underlying storage or the virtual machine. The assumption is that as long as the server restarts (perhaps on alternate hardware), then everything will be “fine.”

However, complex applications don’t always just boot up and self-heal. And it often isn’t until everything is supposedly running again—yet end-users are still complaining—that the application owner is invited to assist in the effort to uncover a solution.

If IT is “all about the app,” shouldn’t that be where the availability plan begins?

<sup>1</sup> Source: ESG Research Report, [Data Protection-as-a-Service Trends](#), September 2013.

<sup>2</sup> Source: ESG Research Report, [Trends for Protecting Highly Virtualized and Private Cloud Environments](#), June 2013.

## Protecting SQL Server Database-centric Applications with DH2i Technology

When IT pros envision the characteristics of a tier-1 application, they often picture a *database*. And tens of thousands of applications today are powered by the database technology known as Microsoft SQL Server. SQL is a relational database-enabling platform supporting all sorts of websites, financial and accounting platforms, engineering environments, shipping and logistics systems, HR platforms, and much, *much* more. If you work with business data, chances are the data is underpinned by a database. And right now, the most prolific and arguably most flexible database platform is SQL Server.

The bottom line is that SQL Server is vitally important to business, but it can be vulnerable without a robust data protection/availability plan. As mentioned, most SQL platforms and their database admins (DBAs) are forced to depend on protection and availability mechanisms that work at the storage or server/VM layer, not at the application layer. As a result, in many cases, data protection operations (other than VSS crash consistency) function almost entirely naively to the SQL-centric applications themselves.

Recognizing this contradiction in circumstance, a company called [DH2i](#) decided to take a different, some might say disruptive, approach to protecting that application data. Recognizing everyone's dependence on data and acknowledging the huge market demand for SQL-powered apps, DH2i has brought data protection and availability capabilities to the SQL platform in a way that seems quite new.

For instance, DH2i's engineers worked out how to smooth away the rough spots between the various SQL versions commonly inhabiting data centers (and there can be many). Regardless of whether those SQL instances are older or newer, or whether they run on physical servers, inside VMware or Microsoft Hyper-V virtual machines, or in the cloud, DH2i software is the abstraction layer sitting on top. Databases and applications built on that abstraction layer become *completely portable*.

### Untying the Hands of DBAs and Application Administrators

With server virtualization, physical metal is abstracted into imaginary metal precisely sized to run whatever OS and upper-layer apps are desired, encapsulating it all into portable form.

DH2i elevated that basic concept two layers higher with its Dx product line of SQL application virtualization software. Instead of a server hypervisor abstracting hardware to create portable VMs, the DH2i DX software abstracts the SQL layer to create portable databases and applications.

As a result, data protection and availability efforts don't need to be tied into generic VM or storage volumes. Instead, an organization's SLAs, application portability efforts, and assured accessibility efforts center on the layer where database apps actually "live."

From a business perspective, this capability is really interesting. It means that if your organization suffers a failure or just has a planned outage, your SQL administrator can bring everything back online faster and more easily (perhaps bolstering his or her own job continuity in the process).

Think about it: Typically, when a database-reliant business app goes down, it remains completely offline until the database supporting it is brought back up. And restoring the database can't happen until the server and storage are functioning again. In other words, until the VM is restored and the storage volumes are remounted, the DBA's and application admin's hands are tied.

It's like taxation without representation. DBAs are liable for application uptime, even though they have no control of the generic underlying infrastructure ensuring that availability. When databases are down, it isn't just "service availability" on the line; the credibility of a DBA who has no infrastructure control is on the line as well.

### DH2i's Solutions for DBAs

However, by abstracting applications and databases at the *database* layer, **DH2i DxHADR** software empowers DBAs and application admins to devise/execute their *own* data protection, high availability, and BC/DR strategies. Any midsize company large enough to have a DBA should appreciate this capability—the DBA will finally get some semblance of

uptime control. Larger enterprises, for their part, should like it, too. They often manage a mix of SQL versions that have proliferated over time, and the related apps may be running on either Hyper-V or VMware.

DxHADR is the latest SQL-savvy product from DH2i, whose portfolio also includes:

- **DH2i DxTransfer**, which enables an IT team to migrate various SQL databases and applications smoothly from one platform to another.
- **DH2i DxConsole**, which enables disparate databases to be monitored, managed, and protected holistically. DxConsole includes the HADR capabilities; it supports shared storage (for consolidation scenarios), and it offers high-availability and disaster recovery capabilities.

The DH2i toolset is designed to reduce application migration-related downtime and improve an organization's availability profile across the board. The DxHADR product is focused specifically on helping organizations meet their demands for assured availability and BC/DR preparedness. Using the DH2i technologies, organizations can improve the disaster recovery and availability capabilities of their platforms to the extent that an administrator could just drag and drop a SQL database from one platform to another, with DxHADR software taking care of everything to reroute and re-instantiate that database. It appears to be an impressive, agile recoverability capability.

## The Bigger Truth

As an IT organization strengthens and matures its infrastructure—transitioning from physical servers to virtualized servers, implementing software-defined data centers, using the cloud, etc.—its databases must become more agile, too. More agile databases translate directly into better application availability for the entire organization, which in turn translates to better end-user productivity and improved profitability.

Application availability efforts must be part of a holistic data protection strategy. To put it more strongly, IT organizations should include availability in the set of data protection activities that they must undertake as they work to ensure access to data. Preserving data, recovering data, *and ensuring the availability of data*: It all needs to happen.

Most application-agnostic approaches to availability have a problem in common. If their automated self-healing methods don't perform as designed and resiliency breaks down, then the application won't be available to end-users because data isn't usable to the application. To make matters worse, it is only after the storage and VMs have failed over but end-users continue complaining that the application owner is brought in to help with the fix.

In IT environments all over the world, the Microsoft SQL Server platform supports and manages business-critical applications and information. Its deployment and distribution truly is pervasive. With this reality in mind, DH2i has taken a fresh approach to solving the availability problem where applications live and breathe, rather than being "application-agnostic" as storage and virtualization-centric methods are. The DH2i approach puts availability control into the hands of DBAs, adhering to the notion that no one knows an app better than its DBA does, so why shouldn't he be at the core of the solution?

SQL DBAs, who have had to depend on application-agnostic resiliency mechanisms for availability, now have a fresh alternative for availability and overall application agility that is certainly worthy of consideration when developing a comprehensive data protection strategy.