



A Dell EMC Integrated Data Protection Appliance solution delivered faster backup and recovery times for databases and virtual machines

Compared to a similarly sized solution from a scale-out vendor, the Dell EMC solution could shorten backup and recovery windows

Your organization likely needs many applications and kinds of data to maintain its overall operational workflow. Interruptions in your infrastructure, whether they're due to a routine backup or a disaster, can affect many users at once and cause a loss of revenue. Minimizing the duration of interruptions allows users to resume accessing data quickly and helps return your operational workflow to a productive pace sooner. Faster data protection infrastructure, such as a Dell EMC Integrated Data Protection Appliance DP5800-based solution, can reduce the time you need to back up and recover business critical virtual machines (VMs) and databases.

At Principled Technologies, we ran several backup and recovery scenarios on two data protection solutions: the Dell EMC Integrated Data Protection Appliance (IDPA) DP5800 and a similarly sized solution from a key competitor (referred to as "Vendor Y" in this report). The IDPA solution backed up and recovered VMs and databases faster than the Vendor Y solution. Our testing featured a mix of backup and recovery scenarios, including a multi-server scenario where we backed up many VMs and databases of different types in a single backup window, something that many large organizations need to do. Faster backups and recoveries can help reduce the risks of losing revenue and violating service level agreements (SLAs), minimize interruptions to critical business operations, and give users more uptime for accessing critical applications and data.



Reduce backup time for production VMs and data

Up to 13% less time to perform a full backup of production-level VMs and databases simultaneously



Minimize downtime during recovery from data loss or corruption

Up to 44% less time to recover production-level databases

How smaller backup and recovery windows help your organization

IT departments set windows of time to back up their most critical applications and data. These windows often happen during low application usage. For companies whose user base extends across multiple time zones, longer backup windows can interfere with application usage and response times.

When it becomes necessary to recover data, every extra minute threatens organizational continuity and increases downtime that can cause lost sales and revenue. If a data set becomes corrupt during normal business hours, for example, IT staff need to restore the most recent data backup so that applications can resume normal activity with the most up-to-date data.

Smaller backup and recovery windows benefit your users and how they work, and thus help your organization. Users depend on data and applications being available to continue their activities and maintain operations. The less time your data protection infrastructure takes to back up data, the better. The less time it takes to recover data, the sooner your employees, partners, clients, and customers can regain access to your data and applications and experience the level of performance to which they're accustomed.



Dell EMC IDPA DP5800 solution

Image provided by Dell

About the IDPA solution we tested

According to Dell EMC, IDPA is a converged appliance that “offers complete backup, replication, recovery, deduplication, instant access and restore, search and analytics, and seamless VMware integration. In addition, IDPA supports cloud readiness with disaster recovery and long-term retention to the cloud, all in a single appliance.”¹

The central component of the Dell EMC solution we tested was the IDPA DP5800, a data protection appliance powered by Intel® Xeon® processors and featuring up to 288 TB of usable active tier storage capacity. The appliance integrates with VMware technology and offers flexibility in the form of many plug-ins and software features that allow organizations to customize backup and recovery capabilities of their data protection environment.

Database administrators using an IDPA solution can also use the Data Domain Boost (DD Boost) plugin, which offloads deduplication and compression to the database server. By using DD Boost, only compressed unique data blocks back up to the IDPA, which can help shorten backups as well as reduce network utilization. In the native application interface, DD Boost could also allow admins to use existing backup procedures after making a single, one-time change to the configuration of the native backup tools.

Reduce backup windows for multi-server environments

If you're running multiple applications and databases, a faster backup process can help maintain operational continuity and mitigate risk. In addition, faster backups allow your organization to set shorter recovery point objectives (RPOs). RPOs are agreed-upon points in time to which IT staff should be able to recover data from a backup (e.g., four hours before an event). Shorter RPOs mean being able to recover more data prior to an event.

In our data center, we simulated seven days of typical operations on a multi-server environment, including data changes and real backup operations on each of the seven days. This schedule of backups follows a typical approach organizations take: conduct full backups once a week and smaller backups throughout the week as data changes. We simultaneously backed up the following mixed workload on both solutions each day:

- An Oracle database and transaction log on a bare-metal server
- A Microsoft SQL Server database and transaction log on a bare-metal server
- 400 VMs running Microsoft Windows Server
- 50 VMs running Linux
- 50 VMs running Microsoft SQL Server

We performed a full backup of the databases and VMs on day 1. On days 2 through 7, we applied updates to the databases and the files on the VMs and then backed up the databases, their transaction logs, and the VMs. For the Vendor Y solution, we completed incremental backups of all assets on days 2 through 7, after completing a full backup on day 1. The Vendor Y solution performs incremental backups by default after completing a full backup. For the IDPA, we completed incremental backups for the VMs and SQL Server database on days 2 through 7, but for the Oracle database, we completed full backups each day. We took advantage of the DD Boost feature to analyze Oracle data at the block level, and compress and deduplicate it on the server before transmitting the data to the IDPA. The IDPA solution backed up the multi-server environment faster than the Vendor Y solution in total and on each of the seven days.

Figure 3 shows that the IDPA solution saved more than 27 hours of backup time for the multi-server environment. That means that an organization with an IDPA solution could see 27 more hours of normal business operations a week compared to an organization with the Vendor Y solution. That's more potential time to conduct operations at a normal pace.

*For brevity, we have rounded all times down to the minute after calculating the differences.

Time to complete a full backup of a multi-server environment




Figure 1: The time in hours and minutes each solution needed to complete a full backup of a multi-server environment. Lower is better. Source: Principled Technologies.

Average time to complete a backup with changes to the SQL Server and Oracle databases




Figure 2: The average time in hours and minutes each solution needed to complete a backup with changes to the SQL Server and Oracle databases. Lower is better. Source: Principled Technologies.

Total time to complete all seven simulated days of multi-server environment backups




Figure 3: The time in hours and minutes each solution needed to complete all seven simulated days of multi-server environment backups. Lower is better. Source: Principled Technologies.

Figures 1 through 3 present different aspects of our findings. Please note, we report times in only hours and minutes for the figures here and on the following pages. The [science behind this report](#) displays times in hours, minutes, and seconds.

Spend less time backing up SQL Server data and transfer less data across your network

In addition to the multi-server environment, we measured how much time each solution needed to back up a single large SQL Server database. We performed a full backup of a 3TB SQL Server database, which we modeled after a data warehouse, on day 1. For days 2 through 7, we changed about 5 percent of the data each day and performed incremental backups for the Vendor Y solution and leveraged DD Boost for Enterprise Applications (DDBEA) to perform full backups for the IDPA solution. We also backed up the database transaction logs on both solutions each day. This backup schedule balances three things: time to back up, storage needed to store backups, and amount of data that could be lost. The IDPA solution completed each of the seven backups faster than the Vendor Y solution. Figures 4 through 6 show different aspects of our timed findings.

Another critical piece of the infrastructure to monitor in the area of data protection is networking. Using fewer network resources for backups can offer organizations room for growth and the ability to use that networking capacity for other things. The more data a backup process pushes across a network, the less the network is available to applications and other processes. As Figure 7 shows, in our single SQL Server database backup scenario, the IDPA solution transferred 32 percent less data across the network than the Vendor Y solution.




Figure 4: The time in hours and minutes each solution needed to complete a full 3TB SQL Server database backup. Lower is better. Source: Principled Technologies.

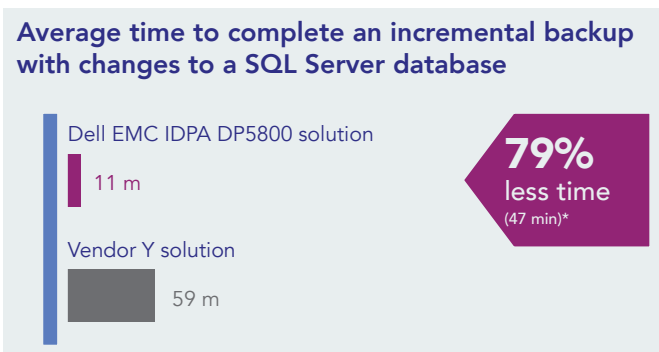


Figure 5: The average time in minutes each solution needed to complete an incremental backup with changes to a SQL Server database. Lower is better. Source: Principled Technologies.



Figure 6: The time in hours and minutes each solution needed to complete all seven SQL Server database backups. Lower is better. Source: Principled Technologies.

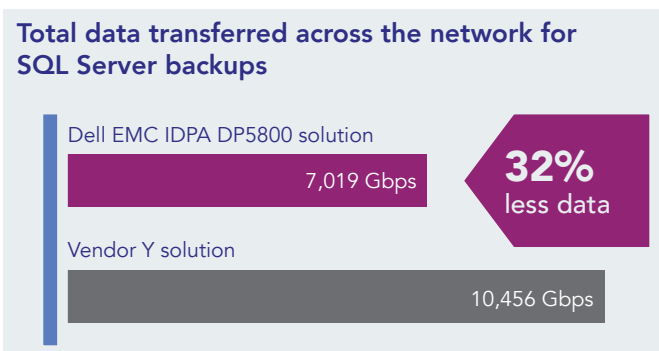


Figure 7: The rate of data in gigabits per second transferred across the network by each solution. Lower is better. Source: Principled Technologies.

*For brevity, we have rounded all times down to the minute after calculating the differences.

Back up an Oracle database in less time

We measured the speed of a single large Oracle database backup for each solution. As in the single SQL Server database scenarios, we performed a full backup of a 5TB Oracle database, which we modeled after a data warehouse, on day 1. For days 2 through 7, we changed about 5 percent of the data each day and performed incremental backups for the Vendor Y solution and used DD Boost to perform full backups for the IDPA solution. We also backed up the database transaction logs on both solutions each day. The IDPA solution completed each of the seven backups faster than the Vendor Y solution. Figures 8 through 10 show different aspects of our findings.

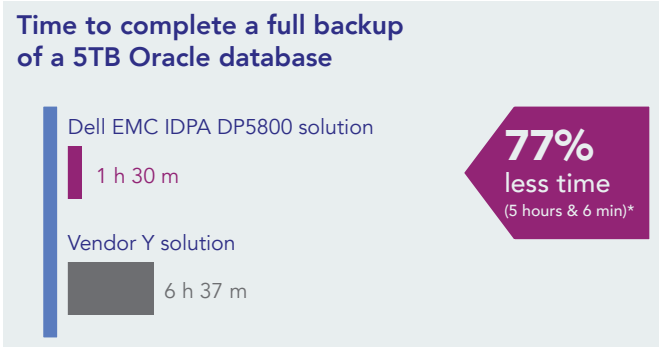


Figure 8: The time in hours and minutes each solution needed to complete a full 5TB Oracle database backup. Lower is better. Source: Principled Technologies.



Figure 9: The average time in hours and minutes each solution needed to complete a backup with changes to an Oracle database. Lower is better. Source: Principled Technologies.



Figure 10: The time in hours and minutes each solution needed to complete all seven Oracle database backups. Lower is better. Source: Principled Technologies.

*For brevity, we have rounded all times down to the minute after calculating the differences.



Shorten recovery time for SQL and Oracle databases

Data corruption, ransomware, or a natural disaster can impede your organization's productivity by taking down applications, destroying data, and prohibiting users from further work. Having a faster data recovery solution in place could allow you to restore affected applications and recover data sooner. Faster data recovery also helps your IT department set shorter recovery time objectives (RTOs), which refers to the contractual amount of time required to recover an application after an event. Shorter RTOs mean faster recovery and less downtime.

To measure the speed of recovery for each solution, we created two common scenarios for recovering data for each of the databases used in the backup testing:

- Full recovery of a database (3 TB for SQL Server and 5 TB for Oracle Database)
- Point-in-time (PIT) database recovery (restoring the database to a specific earlier state)

Figures 11 through 14 show our timed findings.

Figure 14 shows that the IDPA solution saved 10 minutes for the PIT Oracle Database recovery. The default IDPA process provided a more complete recovery because it recovered the full database to a production server (its original location). The Vendor Y solution stops recovery when the database can function from a secondary server that is part of the solution. An organization using the Vendor Y solution can use the data on the secondary server, but they cannot back up updates to the data until the datafiles restore to their primary location. To match the IDPA solution's Oracle backup process, we executed a script to get the Vendor Y solution to restore data to its original location on a production server.

Time to perform full recovery of a 3TB SQL Server database



Figure 11: The time in hours and minutes each solution needed to recover a full 3TB SQL Server database backup. Lower is better. Source: Principled Technologies.

Time to restore a SQL Server database to a specific point in time



Figure 12: The time in hours and minutes each solution needed to restore to a specific point in time for the SQL Server database. Lower is better. Source: Principled Technologies.

Time to perform a full recovery of a 5TB Oracle database to a primary location




Figure 13: The time in hours and minutes each solution needed to recover a full 5TB Oracle database backup. Lower is better. Source: Principled Technologies.

Time to restore an Oracle database to a specific point in time



Figure 14: The time in hours and minutes each solution needed to restore to a specific point in time for the Oracle database. Lower is better. Source: Principled Technologies.

*For brevity, we have rounded all times down to the minute after calculating the differences.



Conclusion

Data protection, in the form of backups and recovery, is vital to the health and the safe operation of your organization. Compared to a competing scale-out solution from Vendor Y, a Dell EMC IDPA DP5800-based data protection solution delivered faster backup and recovery times in several business-critical scenarios. By choosing an IDPA DP5800-based solution for data protection, your organization can set shorter backup and recovery windows, which could minimize disruptions to application performance and user access, help you meet SLAs, and keep your organization's daily operations on track.

-
- 1 Dell EMC, "Dell EMC Integrated Data Protection Appliance," accessed October 28, 2020, <https://www.delltechnologies.com/en-us/data-protection/integrated-data-protection-appliance.htm>.

Read the science behind this report at <http://facts.pt/EEFDscy> ►



Facts matter.®

This project was commissioned by Dell Technologies.

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the science behind this report.