

With data stores doubling in size every 12 to 18 months¹ and consistent increases in virtual server footprints, virtual machine density and throughput requirements², IT teams are understandably concerned about the cost, performance and efficiency of their current data protection infrastructure. When it comes to requisite backup and recovery infrastructure, size matters. The more data there is, the more time and resources are consumed by data protection processes. And as companies look to increase their virtualization footprints and remote office/branch office (ROBO) operations to optimize data centers and expand business reach, they add costly and risky complexity to those processes.

Traditional data deduplication solutions enable IT teams to manage the growth of data stores but do little to conserve data protection resources or reduce cost and management overhead. In contrast, next-generation deduplication provides a single, consistent technology that works across virtual machines to significantly reduce storage requirements and increase performance. Here are 10 reasons federated deduplication makes sense for busy, growing, resource-conscious IT operations.

#1. Reduces IT management overhead with a unified approach to deduplication and backup

Traditional deduplication strategies complicate IT administration by employing multiple, incompatible engines across the IT environment and requiring processes to be run in three locations: the application source, backup server and target device.

Federated deduplication uses a common technology that has been developed for scale, efficiency and performance. And since it's a hardware/software solution, federated deduplication gives IT the flexibility to perform deduplication at locations that work best for the business. For example, a remote office can perform deduplication directly at the application server and back up only new sets of data, while headquarters can do so on a backup server or a target backup appliance.

¹ ESG Research Report, 2012 IT Spending Intentions Survey, January 2012

² ESG Research Report, 2012 Networking Spending Trends, March 2012

A single, flexible solution eliminates conflicts between engines and the need to hydrate data before sending it to a new location.

#2. Maximizes efficiency across virtualized environments.

The downside of otherwise advantageous virtualization is highly redundant data. Deduplication processes identify the new and changed data that needs to be backed up to reduce the overall data volume. Federated deduplication reduces data volume even more than traditional deduplication by comparing much smaller data segments. The higher level of granularity reveals a higher level of commonality, and less data needs to be backed up.

#3. Preserves WAN bandwidth required for high-performance applications and general business operations.

Since federated deduplication generates smaller chunks for backup, it requires minimal network bandwidth. That means business productivity isn't sacrificed to data protection. And even in ROBO environments, IT is able to take advantage of federated deduplication benefits without investing in network upgrades and enhancements.

#4. Increases backup and restore speeds to ensure compliance with service-level agreements and bring new machines online fast.

When supported by federated deduplication, backup processes are three times faster and consume less CPU resources than those that employ traditional deduplication. Data recovery processes are five times faster, which preserves productivity by enabling IT to retrieve lost data and deploy new machines in just minutes. Both performance improvements ensure that data protection will not degrade server, network or application performance.

#5. Improves data recovery effectiveness.

Federated deduplication results in smaller data stores that are easier and faster to search when data restore is required. This gives IT a great opportunity to further improve data restore efficiency with contextual search capabilities. Contextual criteria, rather than exact character matches, increases the probability that one search will retrieve all required data.

#6. Scales easily and cost-effectively with business and data expansion.

Federated deduplication allows operations to grow seamlessly — both in terms of deployment and investment — whenever and wherever

needed. As business locations and virtual machines are added, data stores grow and performance demands increase, IT simply implements new deduplication points. There is no disruption to current operations. "Pay as you grow" licensing enables IT to align investment with actual utilization and eliminate idle resources.

#7. Manages costs by leveraging existing backup resources.

IT teams are able to choose where federated deduplication points are implemented and scale those points as data stores and performance demands increase. Since federated deduplication is hardware and operating system agnostic, IT can use existing equipment when, for example, extending data protection to ROBO locations. And single deduplication stores are shareable across a LAN or SAN, making the federated approach practical even when ROBO locations lack IT staff or expertise.

#8. Streamlines processes and reduces learning curves through centralized management.

Unified, federated deduplication is managed through a single interface. Rather than having to learn and manage disparate processes, applications and user interfaces at a variety of locations, administrators can deploy, manage and monitor backup agents at headquarters and remote or branch locations through a single, easy-to-use interface.

#9. Supports the strategic expansion and dynamic nature of virtual environments.

As virtual machines increase in number, data stores increase in volume and consume more and more equipment, processing, networking and management resources. On top of that, the dynamic nature of virtual machines makes it impossible to maintain real-time telemetry on their location and assignments.

Federated deduplication mitigates those costs with faster processing and smaller data stores that utilize fewer resources. Management burdens are significantly reduced by policy-driven, centralized processes that run consistently across the environment.

#10. Moves deduplicated data anywhere and across software and appliance targets without hydration.

Traditional data deduplication processes that run over legacy appliances generate highly fragmented data. Processes that reconstitute, or hydrate, fragmented data in order to support subsequent deduplication or data

restoration are resource-intensive. That's particularly problematic when processes are consuming bandwidth required by ROBO locations.

Next-generation, federated deduplication eliminates hydration and allows IT administrators to fully protect and restore data across geographically distributed business operations.

To watch demo video of HP Data Protector StoreOnce federated deduplication, visit <http://www.youtube.com/watch?v=nb9RQPQZhgE>

For more information about HP Data Protector, visit <http://www.hp.com/go/dataprotector>

For more information about HP StoreOnce, visit <http://www.hp.com/go/storeonce>