Silver Peak

Duplicating the Success of Data Deduplication

The Benefits of Deploying Local Instance Networking in conjunction with Single Instance Storage

Single Instance Storage (SIS) emerged as one of the hottest storage technologies in the past few years. By identifying redundant data segments and storing only a single instance of information, this technology dramatically reduces storage space and improves network utilization.

However, while Single Instance Storage is extremely effective at improving storage capacity, it was not specifically designed to address the unique issues that arise when disaster recovery is performed across a Wide area Network (WAN). Data reduction certainly provides some benefits on this front, but additional measures are required to overcome network latency and loss, while squeezing even more out of limited WAN bandwidth. In an effort to address all of these requirements, a new WAN acceleration solution called Local Instance Networking (LIN) was born.

Like SIS, LIN uses data reduction techniques to eliminate duplicate data from being processed. However, by applying this concept at the network layer, commonalities can be drawn across multiple applications for increased effectiveness. In addition, this technology is combined with local information delivery and other acceleration techniques to overcome other common WAN challenges, such as latency and loss. When SIS is combined with LIN, enterprises can leverage data reduction to improve storage capacity and accelerate application performance across the WAN. The result is faster, more reliable, and more efficient disaster recovery.

The Emergence of Local Instance Networking

Local Instance Networking works by deploying acceleration appliances in each enterprise location (i.e. on both ends of a WAN connection). The LIN appliances inspect all WAN traffic in real-time and store

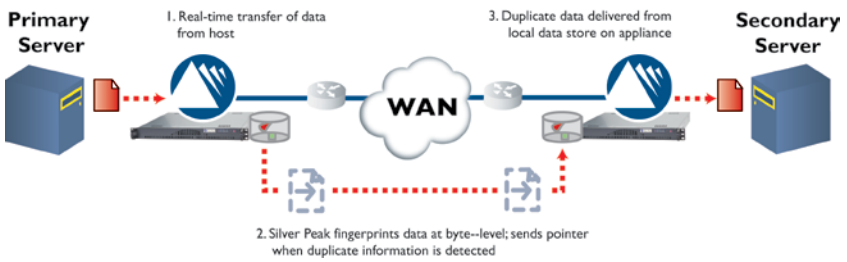

a local instance of information in an application independent data store at the appropriate enterprise location. The local instance is transparently populated based on day-to-day usage, containing a subset of the enterprises working data set that is most relevant to each location. Each instance of information is stored only once per location, enabling an appropriately sized LIN appliance to hold weeks or months worth of data.

LIN appliances examine outbound packets to see if a match exists in the local instance at the destination location. If a match exists, then the repetitive information is not sent across the WAN and instructions are sent to deliver the data locally. If the data has been modified, only the delta is transmitted across the WAN, maximizing bandwidth utilization and application performance.

Local Instance Networking overcomes WAN challenges that often plague common business continuity processes, including backup, replication, and disaster recovery. More specifically, this technology delivers the following benefits:

- **Improve data transfer times.** By delivering repetitive information from local data stores (as opposed to re-sending it across the WAN), WAN transfers are handled at LAN-like speeds. More advanced solutions perform data reduction on both TCP and UDP traffic, delivering significant performance improvements across a wide range of traffic types.
- **Maximize WAN efficiency.** Data reduction can reduce as much as 99% of WAN traffic by eliminating the transfer of duplicate information. When performed at the byte level, repetitive patterns can be detected and eliminated even when the backup/replication solution is performing similar functions at the block level.
- **Increase geographic distances.** By reducing the impact of latency, enterprises can extend the distances between data centers and disaster recovery locations, increasing operational flexibility.

Complementary Solutions

Local Instance Networking is quite complementary to Single Instance Storage. While the latter focuses on improving storage capacity, the former focuses on delivering the best possible performance across the WAN. When LIN is deployed in conjunction with SIS, enterprises typically see a 10-20x performance improvement above and beyond what is achieved with SIS alone. This can be attributed to several factors.

For one, LIN typically provides greater accuracy than SIS when searching for repetitive patterns. This is because individual bytes of data are examined as opposed to blocks, which enables more repetitive patterns to be discovered– even within the same replication stream. In addition, when data deduplication is performed at the network layer, it works across all applications. Therefore, data sent via email, file or web transfer will immediately register as a "hit" when it is sent across the WAN as part of a backup or replication process. In other words, the application itself may not consider the data repetitive, so data deduplication may not work from a SIS standpoint. However, it is duplicate data from a WAN perspective, so LIN would treat it as such.

LIN also works in a bi-directional fashion. In other words, when data is sent from point A to point B, both locations are aware of the information and can deliver the information locally using references, regardless of which direction the traffic is flowing. This can dramatically improve the speed upon which an enterprise can recover data. For example, if information was recently sent across the WAN in one direction as part of normal operations (i.e. replication/backup or simply via email or FTP), then it can be immediately detected when re-sent in the opposite direction as part of the recovery process. Rather than re-transmit an entire data set across the WAN in that scenario, this information can be delivered from local data stores for greater efficiency and performance.

Local Instance Networking minimizes vulnerability within the enterprise and maximizes investments in complementary storage technologies.

LIN appliances are also complementary to SIS because they incorporate other WAN acceleration techniques into the mix for added performance improvement. For example, payload and header compression are often used in conjunction with data reduction to further reduce the amount of WAN bandwidth required for backup and replication.

By providing compression within the acceleration appliance, this functionality can be offloaded from the host replication server, ensuring better scalability and performance. In addition, significant performance improvements can be provided with compression even when non-repetitive information is sent across the WAN. WAN acceleration can also be used to reduce the impact of both packet loss and jitter that occurs when router links are oversubscribed and drop or re-order packets, and they deliver specific enhancements to overcome latency that is inherent to different traffic types, such as TCP. Lastly, Quality of Service (QoS) techniques can be used to prioritize traffic and allocate necessary bandwidth for business critical functions, such as data replication.

Benefits of a combined solution

The most efficient way to save on capacity and improve performance when transferring and storing data is to eliminate redundant information. When data reduction is used as part of Local Instance Networking, the following benefits can be achieved:

- Meet and exceed Recovery Time Objectives (RTO)
- Improve Recovery Point Objectives (RPO)
- Increase geographic distances between data centers
- Avoid costly WAN bandwidth upgrades
- Avoid database synchronization issues that arise when backup and replication tasks are not completed within allocated windows
- Better manage WAN capacity to better handle peak loading, emergency contingencies and business growth

Data reduction is a proven technique for improving the performance, reliability, and efficiency of data backup and recovery. By utilizing this technology in conjunction with other advanced optimization techniques, Local Instance Networking is a strategic IT investment that minimizes vulnerability within the enterprise and maximizes investments in complementary storage technologies.