

Trimming the Branches – Storage Consolidation at the Edge

Analyst: Michael Fisch

Management Summary

Storage consolidation has been an active trend in enterprise data centers for many years. It is the foremost solution for handling the reality of data growth as well as the limitations of direct-attach storage. Rather than discrete storage islands, consolidation separates storage from individual servers and establishes a shared, centralized pool connected usually via a network (i.e., SAN, NAS). This approach is more efficient for similar reasons that carpooling or public transportation is more efficient for commuting – utilization rates are higher and operating costs are lower, which means lower total cost of ownership (TCO). It is also easier to manage and protect data in this architecture.

Why stop at the central data center? **Many enterprises want to extend the benefits of storage consolidation to branch offices, where a significant amount of data also resides.** The status quo of branch office storage is like the Wild West. The deployed systems vary widely and process execution like backup is inconsistent and chaotic. It is not uncommon for branches to rely on office clerks or secretaries to swap out backup tapes and transport them offsite. The result is high costs and inconsistent or poor service levels, which ultimately affect business productivity and profitability.

However, storage consolidation at the edge can turn this scenario around. As in the data center, it can simplify management, improve utilization, and raises overall service levels. **The four key steps to trimming branch storage are:**

- 1) Establish a centralized IT function
- 2) Unify the architecture
- 3) Consolidate branch storage
- 4) Centralize data protection and storage management

There are also three architectural approaches, each with particular advantages, tradeoffs, and technology requirements:

- a) Consolidated, standalone branches
- b) Consolidated, centrally-supported branches
- c) Thin branches

Which is the right approach? It depends on your enterprise's situation and requirements. Read on for more details about why and how to consolidate storage at the edge.

IN THIS ISSUE

➤ Branch Offices – Storage at the Edge...	2
➤ Trim the Branches (Consolidate)	3
➤ Approaches and Technologies.....	4
➤ Benefits to the Business	6
➤ Vendor Considerations	7
➤ Conclusion	7

Branch Offices – Storage at the Edge

The global economy is one big supply chain, labor pool, and market. Naturally, the distribution of corporate offices reflects this reality. Enterprises often have multiple locations, such as a headquarters and one or more branch offices. In the broadest sense, the term *branch office* (branch, for short) refers to retail outlets, sales and marketing offices, call centers, design labs, manufacturing facilities, distribution warehouses, and any other separate location. They can be located across town, around the country, or throughout the world and number in the single digits to the thousands.

Just like headquarters, branches use information technology to support their operations. They have and operate IT systems, though not on the scale of the enterprise's main data center(s). *The table below compares data centers and branch office IT operations.*

In short, data centers and branch office IT operations are alike in that both play an essential role in supporting the business. The difference is that data centers have the focus and resources to deliver a sufficient quality of service, but branch offices frequently do not. They lack resources and, by necessity, are focused on their primary charter (e.g., sales, distribution). IT is a secondary function for them. Many would rather not bother with it, given the choice.

Consequently, branch offices face a number of challenges. The infrastructure is like a smaller version of data centers prior to storage and server consolidation – distributed computing with fragmentation and complexity. Servers with direct-attach storage (DAS) proliferate and create disconnected islands of storage. Utilization is low and management is time-consuming, even more so for the less specialized personnel in branch offices. Routine tasks like upgrades, capacity expansion, and migration cause application downtime. Tape backup jobs may fail or not be performed consistently, risking data loss. System recovery in the event of a disaster, or just restoring a file, is a manual effort that may take hours or days. Distributing data to branch offices - or consolidating data from branch offices - often involves manual, ad hoc methods and the shipping physical media, like CDs. It is a picture of inefficiency.

The result is higher than necessary costs and inadequate or inconsistent service levels. Storage acquisition costs are high due to low utilization, and operating costs are high due to over-consumption of management resources, power, and floor space. Sub-par service levels, in terms of performance, availability, and recoverability, affect worker productivity and increase the risk of data loss and regulatory non-compliance.

The good news is that branch storage

Exhibit 1 - Compare and Contrast

Data Centers

- One or few locations
- Many, larger IT systems
- Numerous, specialized administrators
- Largest budget
- IT is core competency
- Processes are consistent, effective
- Fast data growth, regulatory compliance
- Business depends on data access
- Service levels matter

Branch IT Operations

- Up to thousands of locations
- Fewer, small or mid-sized IT systems
- Limited or no skilled administrators
- Smaller budget
- IT is secondary, support function
- Processes may be inconsistent, less effective
- Fast data growth, regulatory compliance
- Business depends on data access
- Service levels matter

presents a ripe opportunity for improving IT efficiency. Some might call it low-hanging fruit.

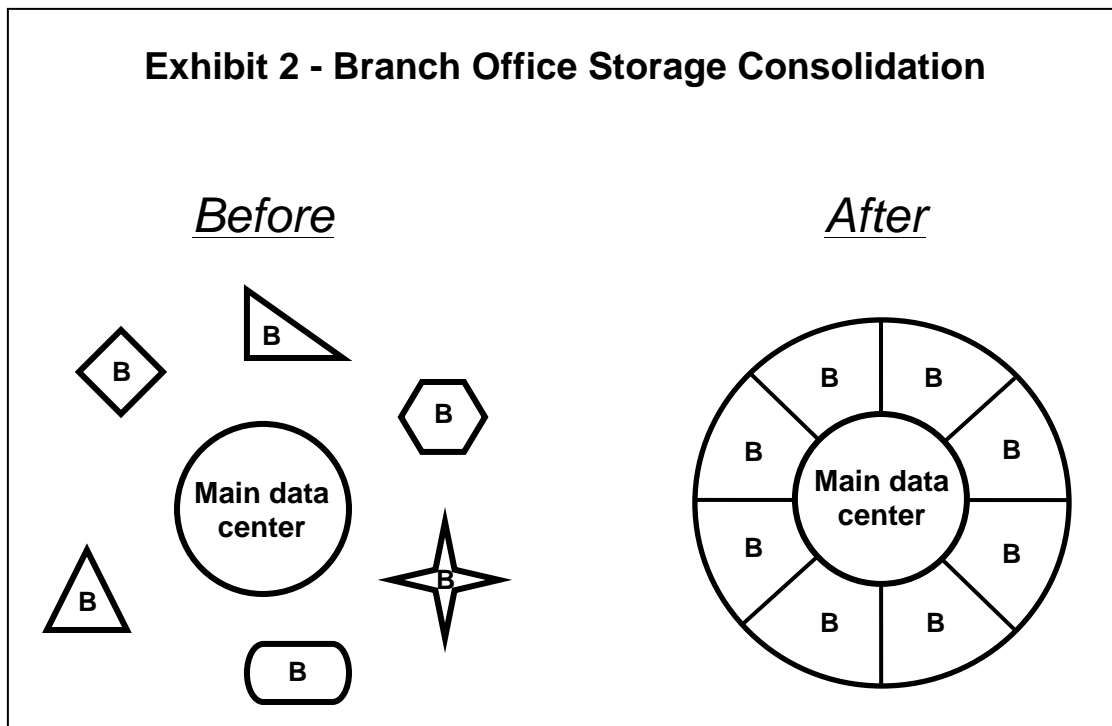
Trim the Branches (Consolidate)

Like trimming the branches of an overgrown tree, consolidating IT operations is the answer to the challenges described above. It allows companies to bring their efficiency in line with the data center operations and safeguard valuable information assets. We recommend the following steps.

- **Establish a centralized IT function** – If it does not already, the central IT department should have primary responsibility for branch IT operations. This first step is important. If branch offices remain in control over local storage procurement and management, then the helter-skelter situation described above will not change. Only a centralized approach affords the opportunity for broad-scale efficiencies. Executives will need to support this change, since it is a political and cultural issue as much as anything.
- **Unify the architecture** – The technologies and processes used for branch offices should have a common design, which provides economies of scale in procurement, testing,

integration, deployment, and management. The architecture should meet most all or most branch office requirements, so the need for local customization is minimized, as that would add complexity and cost.

- **Consolidate branch storage** – Storage consolidation has been an active trend for quite a while, and much has been written about it already. The premise is that it is better to have a single, shared pool of storage than many disconnected islands. Utilization is higher; management is simpler; service levels are better; and the infrastructure is more adaptable to changing business requirements. Consolidation can be physical (i.e., bigger boxes) and/or logical (i.e., unifying multiple physical assets with virtualization). Beyond local consolidation, enterprises have the opportunity to create synergies between branch offices and the main data center using wide area network (WAN) links to centralize some equipment, data, and processes. (*See Approaches and Technologies on the next page.*)
- **Centralize data protection and storage management** – This means remote management from the main data center by the central IT staff. As much as technically possible, centralized management and auto-



mation of branch office storage is better than local efforts. It requires fewer administrators and generates fewer errors.

The graphic on the previous page loosely depicts concept of branch storage consolidation. The *before* picture is chaotic and disorganized, with no connection between the main data center and branch offices. The *after* picture is shows a consistent branch office architecture and closer integration with the main data center. The degree of integration depends on the approach taken for consolidation and centralization.

Approaches and Technologies

There is more than one way to trim a branch. In fact, there are three architectural approaches to branch storage consolidation, each with particular advantages and sets of supporting technologies. It is incumbent upon enterprises to decide which is most suitable for their requirements.

(1) *Consolidated, Standalone Branches*

Consolidated, standalone branches operate independently. This approach consolidates storage and centralizes data protection and storage management at each branch, creating many smaller versions of a consolidated data center. The branches continue to support their own applications, servers, and storage equipment, but more efficiently than before. Each branch needs at least one IT administrator. The role of the central IT department is to supply the architecture, management processes, training, and secondary support. Several technologies support this approach.

- **Networked storage (SAN/NAS)** – Separates storage from servers, consolidates it on a network, and makes it broadly accessible to multiple applications and users. A SAN provides access to a pool of block storage for application hosting, whether via *Fibre Channel* (best performance) or *IP* using the *iSCSI* protocol (simpler, lowest cost). NAS provides shared file system access over an IP network using standard *CIFS (Windows)* and *NFS (UNIX, Linux)* protocols. Some products offer both NAS and *iSCSI* SAN in one, which is a nice capability for branches. For very small environments with one or few collocated

servers, a shared DAS array may be acceptable.

- **Centralized backup and restore** – Backs up data on all servers, desktops, and laptops at a branch from a single management console. It should be easy to use, automated, and support the applications and operating systems in the branches. Furthermore, using disk as a backup target instead of tape is faster, more reliable, and avoids handling physical media. Electronic vaulting sends copies offsite over a WAN for disaster recovery purposes, again, without handling tapes. Solutions that integrate with point-in-time (PIT) copy technology enable backups to be performed non-disruptively.
- **Centralized storage and SAN management** – Manages storage and/or SAN assets from a single console, even heterogeneous environments if that is what branches have. A SAN requires a multiplicity of physical and logical components to interface and work together, such as files systems, volumes, host bus adapters, cables, switches, routers, storage arrays, RAID groups, and LUNs. Using individual device managers to coordinate the overall environment is difficult. However, SAN management software simplifies the end-to-end process with monitoring, mapping, configuring, and automation.
- **Storage resource management (SRM)** – Gives a mountaintop view of data in a branch. “*What data is out there and where is it?*” SRM software scans file systems and databases, tracks data by type and amount, and often correlates it to specific storage devices. Branch IT staff can use this information to perform a variety of tasks, like capacity planning and expansion, archiving inactive data, load balancing, and quota management.

The benefits of consolidated, standalone branches are flexible, robust local IT operations and improved economics over fragmented DAS environments. Efficiencies come from using one architecture across the branches and consolidation within each one. It is best for branches with many users, demanding performance requirements, and multiple, possibly unique applications. This approach gives

branches the most local control and flexibility, but leaves some efficiencies on the table since it only partially leverages the central IT department.

(2) Consolidated, Centrally-Supported Branches

Consolidated, centrally-supported branches leverage the central IT department to a greater degree. This approach consolidates storage at the branches and shifts the responsibility for storage management and data protection to the main data center. Though application servers and storage equipment still reside at the branches, the central IT staff monitors and manages them using remote tools (except for tasks that require touching the devices). They also replicate data from branches to the main data center for backup and archiving. This approach depends on a wide area network (WAN) link to each branch. The central IT department supplies primary management support, instead of secondary, as well as the technology architecture, processes, and training. Branches have fewer IT responsibilities and need less expertise. Three technologies support this approach.

- **Networked storage (SAN/NAS)** – Same as above.
- **Consolidated backup and restore** – Gathers incremental backup data at the branches (i.e., changed since last backup) and replicates it to the main data center. The backup data is then integrated into the central data protection and archiving processes and systems. Restores flow in the reverse direction. Asynchronous remote mirroring is another way to replicate data from branches to the main data center.
- **Centralized, remote storage and SAN management and SRM** – Functions like the tools described above but also can be managed remotely.

The benefits of consolidated, centrally-supported branches are robust local service levels and greater economics than consolidated storage alone. Efficiencies come from a single architecture across branches, local storage consolidation, leveraging central IT staff and data protection systems in the main data center, and reducing IT staff at the branches. In this way, enterprises increase the storage capacity

managed per administrator and lower management costs, the largest component of storage TCO. Better service levels come from more consistent and automated management processes, especially data protection and recoverability. It also affords the opportunity to archive corporate data from the branches at a central location – for regulatory compliance and corporate governance, among other reasons. In exchange for greater efficiency and fewer IT headaches, branches cede some local control and flexibility.

This approach strikes a sensible balance between service levels and cost. Data access and transaction response time are fast since application servers are on the local LAN. Only management tasks and data protection rely on WAN links. Technologies for implementing it are established and readily available.

(3) Thin Branches

Thin branches are the ultimate in efficient IT operations. This approach shifts not only storage management and data protection to the main data center, but also data, storage, servers, and applications. In other words, all business applications run in the main data center and branch users access them over WAN links using remote access and caching technologies. Branches contain no storage except PC drives and possibly a caching appliance, so they do not need IT staff for storage. The central IT department owns and manages all storage equipment and processes in the main data center. Three technologies support this approach.

- **Remote access platforms (also called terminal services)** – Allows users to remotely access applications running in the main data center. The user's session and data remain in the data center, and the remote access platform sends presentation screens over a WAN to the client PC or other device. It requires relatively high bandwidth connections, and user response time depends on bandwidth, congestion, and latency. It is not as fast as LAN access. Interactive or "chatty" applications, like Microsoft *Office* especially, can affect performance. Business requirements will determine whether this technology can achieve acceptable performance for branch users.

- **Proxy servers for content and application delivery** – These appliances cache frequently- or likely-to-be-accessed data in branches, which can include file and some application data. Users experience LAN response time for cached data and WAN response times for data fetched from storage in the main data center. There can be a coherency risk if the source data changes before the cache is updated.
- **Wide Area File Services (WAFS)** – Provides branch users read/write access to files hosted at the main data center at near-LAN speeds. How “near” it is depends on the particular product’s features as well as WAN latency, application protocols, and file size. WAFS solutions employ one or more of several technologies to shrink the gap between data center and branch: data reduction (indexing algorithms), compression algorithms, TCP optimization, caching, and file system or application protocol acceleration. They usually require appliances at both ends, but not always. Though the technology is in the earlier stages of adoption and maturity, there is already a variety of products in the market. Moreover, WAFS is actually part of a broader category that includes WAN optimization and application protocol acceleration, which some refer to as Wide area Data Services (WDS).

Thin branches maximize storage efficiency and cost effectiveness. Data and storage is consolidated in the main data center, where the central IT staff manages and protects it. Branch data occupies a slice of the consolidated storage infrastructure for all enterprise information. This approach affords the greatest degree of consolidation and centralization. Performance in branches is at best near-LAN, since users are dependent on the WAN link for all data access, except for anything cached locally. If the WAN is congested or offline, then users must wait until it is restored.

Factors Affecting Consolidation

Of course, these three approaches are archetypes. An enterprise might choose to implement a modified or in-between version. For instance, it is possible to host some applications in the main data center for remote access and others in the branches for local access. Factors that help determine the right

approach for an enterprise are:

- Number and locations of branches
- Number of users per branch
- Applications
- Business requirements for performance, availability, recoverability, and security
- Criticality of data
- Cost and availability of WAN links

Benefits to the Business

The business benefits of branch storage consolidation include the following.

- **Lower acquisition costs** – A consolidated pool of storage allows capacity to be readily shared and reallocated among servers, minimizing the amount of unused or “overhead” capacity and improving utilization. It also allows servers and storage to scale separately, so upgrading storage does not require a server upgrade, and vice-versa. Less overhead also reduces power, cooling, and floor space expenditures, a component of operating costs. The amount of savings depends on whether the storage is consolidated within each branch (good) or to the main data center (best).
- **Lower storage operating costs through simplified management** – This is a major benefit because the cost of managing storage over its useful life can be several-to-many times its acquisition cost. With consolidated, networked storage, common tasks like scaling capacity, migrating data, and performing backups are faster and easier, especially if using centralized management tools. This empowers each administrator to handle a much more capacity. Once again, the degree of savings depends on whether consolidation is in each branch or to the data center.
- **Enhance business productivity** – Storage is foundational to all IT applications that, in turn, support business processes and broader enterprise objectives, both in headquarters and branches. By delivering an improved quality of storage service, consolidated storage can have a positive effect all the way up the business chain, ultimately impacting productivity. But also keep in mind that consolidating branch storage to the main

data center (*approach #3 above*) would somewhat reduce performance and availability delivered to the branches. This is the question each enterprise must answer: *Are the extra savings worth the impact on service levels?*

Vendor Considerations

When considering branch storage consolidation, you will want to find a vendor who can take the right approach and build the right solution for your enterprise. Below are some vendor characteristics to consider.

- **A storage product line with breadth and scale to meet both branch office and data center requirements** – The product line should have broad capabilities in terms of SAN, NAS, and centralized storage and data management software. It should also have right-sized solutions for small branch offices to large data centers - in other words, not too big for the edge or too small for the center.
- **SAN and NAS solutions that start small in capacity and price and scale up in a non-disruptive, pay-as-you-grow fashion** – This is important for branches that have small-scale storage needs today but could grow much larger tomorrow. Scaling up should be smooth; avoid rip-and-replace scenarios.
- **Enough performance and availability** – These service level characteristics are a function of the storage platforms themselves as well as storage and data management software layered on top.
- **Ability to deliver a complete, integrated solution** – More than pieces and parts, this means integrated solutions, professional services, and partnerships with other vendors. Branch consolidation projects are much more than point solutions. They span the storage infrastructure as well as servers and applications. A vendor's products should interoperate with third-party products to form complete and effective solutions.
- **A geographic presence that matches your enterprise** – Branch offices are by nature spread out. Therefore, a global enterprise needs a vendor that can deploy and support solutions around the world.

- **Reliability and a proven track record** – Investing in a unified architecture across an enterprise implies a lot of trust in the chosen vendor. A successful history in networked storage and software is a positive indicator of a vendor's ability to be a long-term storage partner, as is ongoing investment in R&D and innovation.
- **Easy implementation and management** – If branches will deploy the storage solutions themselves, then ease of installation is important. Pre-integration and deployment wizards are helpful features. If branches will also manage them, intuitive, easy-to-use management software is also important.
- **Reasonable pricing** – You get what you pay for, of course, but the value should be there.

Conclusion

If the storage in your enterprise's branch offices looks more like the Wild West than the civilized world, it is time to start trimming. Branch storage consolidation can yield significant cost savings, better service levels, and productivity improvements. It is just good business.

The operative words are *unify*, *consolidate*, and *centralize*, but the right approach depends on your enterprise's particular situation. **So, appraise your requirements, choose an effective approach, and work with a vendor that can help take you there.**



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over ten years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and "clipper.com" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "Navigating Information Technology Horizons", and "teraproductivity" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.