THE CLIPPER GROUP
Observer™

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2004045 May 18, 2004

Curbing Chaos with Bandwidth Management — A Tale of the English Countryside

Analyst: Anne MacFarland

Management Summary

The imperative was simple. Rural schools in a sub-region of England (specifically, Southern Holderness, just east of Hull) are underperforming. The government is funding a project called the Rural Achievement Zone (RAZ), to change this situation and ICT (Information Communications and Technology) will be a large part of it. The goal: to build a system that will revitalize the education process, and make it useful to the adults of the community as well. There is no ROI, just improved educational results and motivation. There is no TCO to be calculated, but the budget is meager. You will have to justify your expenditures against educational targets. Success – and failure – will be obvious. Best of luck, of course.

Ray Ford, Director ICT and E-Learning of the Rural Achievement Zone, faced this challenge when he was asked to set up the ICT strand of the RAZ at 13 primary schools and one secondary school to provide 21st century learning solutions for a cluster of rural communities. At age 5, each child gets a network login that will be kept until each is 18. Starting around age 12, they are transitioned from "educational" software to more standard productivity software (Microsoft, Adobe, Macromedia, etc.) that, together with Internet access, are used in everyday lessons. Teachers can customize functionality for students who need it. Adults use the system for education and training. In response, Ford designed **an open system infrastructure with virtualized SAN storage, extensive remote management and automation, and a Sandial switch to manage the chaos by managing bandwidth.**

What chaos? To start with, network utilization is sporadic. School days are segmented by the hour, so everybody is logging on and off simultaneously and changing his or her profiles and home directories constantly. Because the system includes e-mail, it is also used seven days a week. Of course, there are a few student hackers trying to bring the system down, but that is normal in this environment. These miscreants, as part of the community, should be deterred, but, if not, must be tolerated – and the effects of their efforts must be contained. Restoring files is a fact of life, not a dreaded event. When a restore is needed, the required bandwidth must be available to complete the task expeditiously.

The equipment is distributed among the schools, as there is no money for purpose-built data centers. The region cannot afford to upgrade to dedicated electrical systems, so the system must have the resilience to operate under normal school system's environmental conditions (temperamental, at best) and use locally-generated (UPS) power for up to eight hours, as is needed. All schools are also open weekday evenings for adult education and training. Students and adults also have network access to a sound and video production studio and expertise, due to Ford's astute choice of colleagues. This is very popular – and is a real bandwidth hog.

This system has become an integral part of village life. Students regularly come to school early and stay late, and adult attendance in the evening is high. Due to the success of the initiative, it is being expanded to two other areas each 50 km away. Ford's optimism sees this challenge as an opportunity to build in still more system resilience. There is much to be learned from Ford's experience. For more details, read on.

IN THIS ISSUE

- The Design and Its Constraints 2
- Operations 3
- What Didn't Work 3
- Expansion Plans 3
- Three Take-Aways 4
- Conclusion 4

The Design and Its Constraints

Ray Ford started small, wheedling for money and buying used equipment. As the results of the program became evident, money became easier to find, but the budget was still not enough to pay business prices for the enterprise-class functionality that the situation demands. Instead, Ford has saved money where possible and built in automation and redundancy, so that the system can be managed by a staff of three.

Ford has designed an infrastructure that is to last five to fifteen years. It comprises 70 servers, 700 PCs, and about 5000 users. The servers are older HP and Compaq boxes, some as large as 4-way systems, though the project is standardizing on 1U pizza boxes with 4GB RAM and dual-port Emulex host bus adapters. Two years ago, the SAN had 1 GB connectivity and dual *Vixel 7200* switches. It now runs at 2 GB/s, and is managed by an internally-redundant, next-generation Sandial director-class switch.

Every school has a 2MB connection to a main secondary hub. Connectivity between the main secondary hubs and the outside world is being upgraded from 4 MB to 30MB. Using Microsoft's *Active Directory*, users have a roaming profile, so they can log on from anywhere and have access to their network resources as if at school.

About 50-70 GB of data is kept at each site. A file-level incremental copy is done daily to an off-site staging disk array for DR (disaster recovery) and instant restore. A logical snapshot then goes to a back-up server, which replicates the data to virtual tape (a SATA array) and then to LTO tape. The multiple recovery options of this structure are critical to the resilience that this environment needs.

Storage Virtualization

All storage has been virtualized for the last three years. Ford has chosen to go with Fibre Channel (FC) storage arrays of various vintages and brands, and plans to use a secondary tier of cheap SATA arrays for data with lower performance requirements. Enough redundancy is built into the system to make the use of SATA a safe as well as economical way to go.

Educational ICT environments are combat zones that demand extreme flexibility and quick response to system changes. Multiple clustering solutions are used, frequent hardware upgrades are the norm, and migrating volumes from one array to another is commonplace. These constant changes make traditional LUN masking an impossibly cumbersome approach. Instead, agents of the StoreAge Networking Technologies' virtualization appliance running *StoreAge Volume Manager* and other tools manage the LUNs as well as dual-pathing and failover. Large LUNs (1 terabyte, minimum) are given to the StoreAge virtualization appliance. This allows creation and presentation of any size volume to any host on any OS through a browser (from anywhere) and, more importantly, the removal and reusing of storage capacity that can present any size volume to any host. As Ford says, "This sector needs to use everything as flexibly as possible with no waste. Using virtualization allows us to work with any array, any switch, any server, and any operating system. The result – vendor neutrality and no lock-in."

Bandwidth Management.

Ford quickly realized that the core virtualization technologies used to maintain the StoreAge system (MultiMirror[1], MultiCopy[2], MultiView[3] and serverless back-up[4], required very precise bandwidth control within the SAN, otherwise these system maintenance workloads could interfere with access speeds to the data that was mission critical. Whether by local power problems or by hackers, some nodes go down, and recovery has to be a full-bandwidth effort. **Being able to dynamically manage the bandwidth given to different applications was the only way to make this system work efficiently. At the time RAZ started, no storage switches had the capability to manage bandwidth. The search was on.**

Two years ago, Sandial Systems (of Portsmouth, NH) had not yet released its *Shadow* switch[5] – but Ford heard about it, and was able to get one. Sandial saw this as a test-bed for what was possible with their product. Ford saw this as exactly what he heeded.

Ford and his colleagues had the skills to use the next generation and the Fibre Channel capabilities of Sandial's *Shadow* switch. RAZ is using Fibre Channel over IP (FCIP) and is looking at implementing iSCSI, where appropriate. Because using redundant switches was never going to be a budget option, Ford uses left-hand/right-hand failover functionality within the switch.

Ford feels that the combination of storage virtualization and commodity (as opposed to high-end) servers with Sandial's bandwidth management can do just about anything. Moving data from one site to another non-disruptively using a

---

[1] Local and remote mirroring of data

[2] Physical copying of data between arrays

[3] Point-in-time snapshot use by multiple hosts for testing, back-up, etc.

[4] External data movers

[5] See **The Clipper Group Navigator** dated February 18, 2004, entitled *Sandial Shadow 14000 - Optimizing the N in SAN*, available at http://www.clipper.com/research/TCG200411.pdf.

point-in-time snapshot makes critical use of Sandial's bandwidth management. Serverless backup, critical in a distributed environment, needs to be throttled down during the day, particularly at the top of the hour when user turnover is high, and ramped up at other times, using bandwidth settings. Moving data between arrays to provision storage where required without application downtime is also critical and can be accomplished most expeditiously with bandwidth management.

High Functionality Software

In addition to Sandial and StoreAge mentioned above, another key partner has been Computer Associates, whose educational pricing and academic partner program have allowed RAZ to get the enterprise features that they need and excellent support at a price they can afford. RAZ is moving from a less-scalable web portal to CA's *CleverPath* portal to present content to all of its various users. Ford also uses CA's *BrightStor* serverless backup, which he praises for its ability to work well with virtualized storage and with data movers from Crossroads Systems, another helpful partner.

The availability of the high-availability features that the inherently hostile educational environment needs are not always a part of the bundles offered by vendors at "educational" prices. Ford feels strongly that there is an urgent need for a better way for educational institutions to get the features they need at prices they can afford. Education is a market separate from the SMB-to-large enterprise continuum – but it is one that can bake brand loyalty into every child who grows up using an invisibly-persistent infrastructure.

Operations

Teachers simply use the RAZ system as a learning space. It has allowed them to become tutors, who can ask for customization for individual students as needed. The entire RAZ infrastructure is remotely managed and all functionality is redundant, so the system stays up until any needed repairs are made during the overnight hiatus. This seems to work much better than educational technology systems where some of the teachers double as part-time network administrators. With systems of this complexity, Ford feels that decentralizing the responsibility for operations and management is not viable or recommended.

Ford freely admits that the end users are his worst enemy. Young children break the system by natural curiosity and accidental actions, but this is expected. There are, however, always half of a percent who find joy in destruction, and by the time they are twelve or thirteen, they are really good at it. Without the time or inclination to track down and punish the hackers, Ford has instead built a secure system that can track, isolate, and record their attempts, and provide teachers with the evidence to challenge this behavior directly with students and parents. This has worked well.

What Didn't Work

The first Virtual Learning Environment (VLE) using the RAZ system to administer classroom tests did not work well. The databases needed to support such efforts were a complexity that there is no staff to support. The interactivity of the testing process used more bandwidth than RAZ could afford. The old way of locally-administered testing was effective, so this initiative was dropped.

Firewalls and caching engines were initially installed at the core site only. The caching, while greatly improving data access speed, provided only very basic filtering. Due to the extreme granularity of filtering and customization the K-to-12 educational environment demands, an application plug-in approach worked better. In the student population of 5000, there are not just many *classes* of users but, in many cases, individual instances with particular needs.

Multicasting of video required switching capabilities that RAZ could not initially afford. Ford plans to use content distribution, edge caching, and pre-positioning management to provide multicasting capability when the demand requires it.

With the availability of Windows 2003, RAZ was able to move from thick servers to high-availability clusters. Ford is also putting more and more applications functionality through a browser, and moving away from location-specific applications. He feels it is not only more efficient, but also makes the content available 24x7 and provides users with a single point of access to all resources.

Expansion Plans

RAZ is working to expand interactive content beyond the existing e-mail and Web browsing, using a dynamic HTML front end that can recognize users and present appropriate material to them. With users ranging from age 5 to 80, there is a wide range of granular content filtering that must be supported. Ford feels that the stateless quality of HTML will allow him to meet these needs for granularity without unnecessary infrastructure expansion.

RAZ itself is continuing to expand. Two more educationally-challenged rural areas have been targeted in Bridlington and Goole. This summer the system will triple in capacity to support these areas.

This will enable a 4-domain SAN design between the primary site in Withernsea, the new sites on Bridlington and Goole, and the disaster recovery mirror site at a primary school in Keyingham. Each site will have four-to-six terabytes of virtualized storage initially. These synchronized domains give resilience to the environment.

The environment will remain administratively united. Ford sees federated management of the environment as inherently problematic. While he would like to get more sleep, the efficiencies of providing access to the same applications with a common interface are highly appreciated, and the local administration and customization let schools take ownership of the ICT capabilities.

Ford is proud of his system. It has transformed how teachers use ICT and what users of all ages can do in an educational setting. Reports become videos. Users can capture content on their cell phones and paste it into their creations. This appeals to more than just children, and can transform this back-end-of-beyond into an incubator of creativity.

Three Take-Aways

(1) The Importance of Resiliency and Automation

Life is imperfect. The era of building for perfection has been succeeded by a more pragmatic need to build a system that can quickly recover from an unexpected array of unanticipated events, and that can recover over time from serious disaster. Integrators must build this resiliency into the architecture, using automation to cover the *what-if*s, rather than an eternal expensive vigilance. This approach will let administrators concentrate on their customers, the end users.

(2) The Passive-Aggressive Nature of Granular Needs

Johnny's sudden need for an advanced Adobe functionality and the filtering to curb Maggie's suddenly-precocious curiosity puts those users in the driver's seat, and can leave administrator fumbling for their own seats and mumbling imprecations, if storage and applications are manually managed separately. It is not just that the role of the system as the center of a Ptolemaic universe will not serve, but that traditional elements of that system may have to be rethought. Building for resilience differs from building simply for scalable performance. Tight integration can be counterproductive when change is rampant. If firewalls by their nature encumber, more-intelligent intrusion detection and other forms of security, must be developed to support a diverse user base. There is no one size fits all.

In a customer-centric environment such as RAZ, **granular needs trump traditional ideas of efficiency and architectural purity.** RAZ's range of needs may seem unusual, but, if users had their way, such needs for customization would not be uncommon. **Use what is needed to get the results required. Build to suit, not to style**.

(3) The Value of separating the Physical from the Logical, and the Consequent Need to be Flexible, and Manage Granular Needs by Bandwidth

Storage virtualization is crucial to the effectiveness of the RAZ system, because it enables the provisioning of storage resources on an *as-needed* basis. Equally crucial is bandwidth management. Both functionalities make the high availability hardware solutions, like cluster failover, that a hostile environment demands cost effective. In some other environments, processor virtualization, and the environment-virtualization of grids, might also be necessary. **Virtualization, and the tools to manage virtualized environment, should be familiar in any environment that must cope with rapidly-changing demands.**

With the linking of devices on a common network and the use of virtualization comes a move of the point of control from the device to the service. We see in the example of RAZ that events will change what levels of service are delivered to what applications. When restoring environments is necessary, surplus bandwidth must be drained from other tasks – but not to the point of making the system inoperative. **The intelligence to manage the bandwidth must come from an intelligent switch.**

And so, with the example of RAZ, we have **a service-oriented architecture, built initially out of second-hand components and later out of next generation hardware and software, that gets the job done.** The magic bullets may be virtualization and bandwidth management, but the real magic is in minds of the people who design, build, operate, and evolve the system.

Conclusion

This is the story of how technology can enhance the educational process. It is a story of how technology can enhance life in small, rural communities. But it is also a story of how to build a rugged, flexible environment, tailored to daunting parameters, while containing costs – a universal story that has much to teach us all.

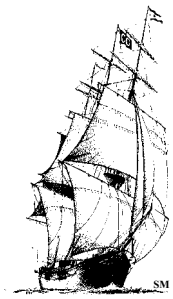

About The Clipper Group, Inc.

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

Anne MacFarland **is Director of Enterprise Architectures and Infrastructure Solutions for The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 28. (Please dial "1-28" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

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.