Implementing A SAN

Rob Strechay

You knew about the speed vs. security trade-offs on LANs and WANs, but did you know that SANs have trade-offs, too? Think speed, port density and storage capacity.

The cost of acquiring and implementing a storage area network (SAN) has recently started to come down to earth, but not because SANs are becoming Ethernet-like commodities. Instead, like most publicly-held tech companies, many SAN vendors are trying to cut their inventories. The lean times are also putting consolidation pressures on smaller storage start-ups, forcing some of them out of business or into the arms of the larger, more-established competitors in the field.

So this can be a good time to make a deal on SAN infrastructure, or on direct-attached storage (DAS) or network-attached storage (NAS) gear—assuming you have done your homework and know what you need. But don't shortcut the fact-finding and business-case processes. Every storage implementation has unique requirements and must be custom fit to the existing applications, servers, disk arrays and management software: That's another reason that SAN, DAS and NAS solutions are not yet commodities.

Recently I took part in designing and implementing the SAN system shown in Figure 1. It took about five months, but we had gone through requirements gathering beforehand, and we understood what we were looking to achieve from a business perspective (see "Getting Started With SANs" pp. 26). You may be able to do the design faster, especially if you have a less complicated set of requirements, but everyone should allow for extra time, money and skills in the implementation phase.

Rob Strechay is the principal consultant with Strechay Communications and the former manager of network operations for Wheelhouse Corp. He can be reached at rob@strechay.com. This article is based on his recent experience with Wheelhouse Corp.

Why We Needed A SAN

As an outsourcing vendor supplying marketing platforms and equipment on a contract basis, our business goal was to alter our storage systems to improve service, as well as to gain flexibility and control costs. We knew we could improve our service and reduce costs by using the same storage infrastructure for both Windows- and Unix-based projects, thereby avoiding vendor-specific storage purchases and speeding up the development time.

We also wanted to more flexibly leverage our storage assets to better serve clients. Most of our work involved developing and hosting large customer relationship management (CRM) systems, and some of the staging systems we used to build these CRM systems had changing requirements. The staging servers could have multiple loads occurring daily, and their disk arrays, consisting of 18- to 72-gigabyte (GB) disk drives, would be allocated for the duration of the system's construction. Once the system was constructed, the space would be reallocated to a different project. Projects typically lasted three to six months.

On the ongoing CRM hosting side, some clients had service level agreements (SLAs) that directly stated the amount of downtime permitted for their CRM updates. Some had large updates that ran on the weekends, while others were updated nightly. The updates could take three to 26 hours, while the volumes ranged from 1 GB nightly to more than 5 GB weekly.

We also had application and other types of servers that sometimes needed additional storage space for short time spans. Prior to the SAN project, we were running direct-attached storage in various RAID configurations (from RAID 5 to RAID 0+1) on every server. The performance on the data loading of the CRM systems was becoming an issue; some of the updates already were experiencing more downtime than their contract allowed and others were getting close.

From an administration standpoint, managing separate, fluctuating storage requirements for each of the various corporate, staging and product servers also was becoming an issue. We wanted to keep our existing volume management software, because the staff was familiar with it, but we often had to physically take drives from one RAID group on one server, rebuild and redeploy them to support another server. We worried about our many single points of failure. This is common in most DAS shops, but our business goals made us want to minimize our exposure to this type of issue.

Finally, we had a tight timetable: three months for selection and two months for implementation.

Who We Invited, What We Wanted

We made a list of vendors for consideration, based on our team's knowledge and past relationships,

FIGURE 1 Current SAN Architecture

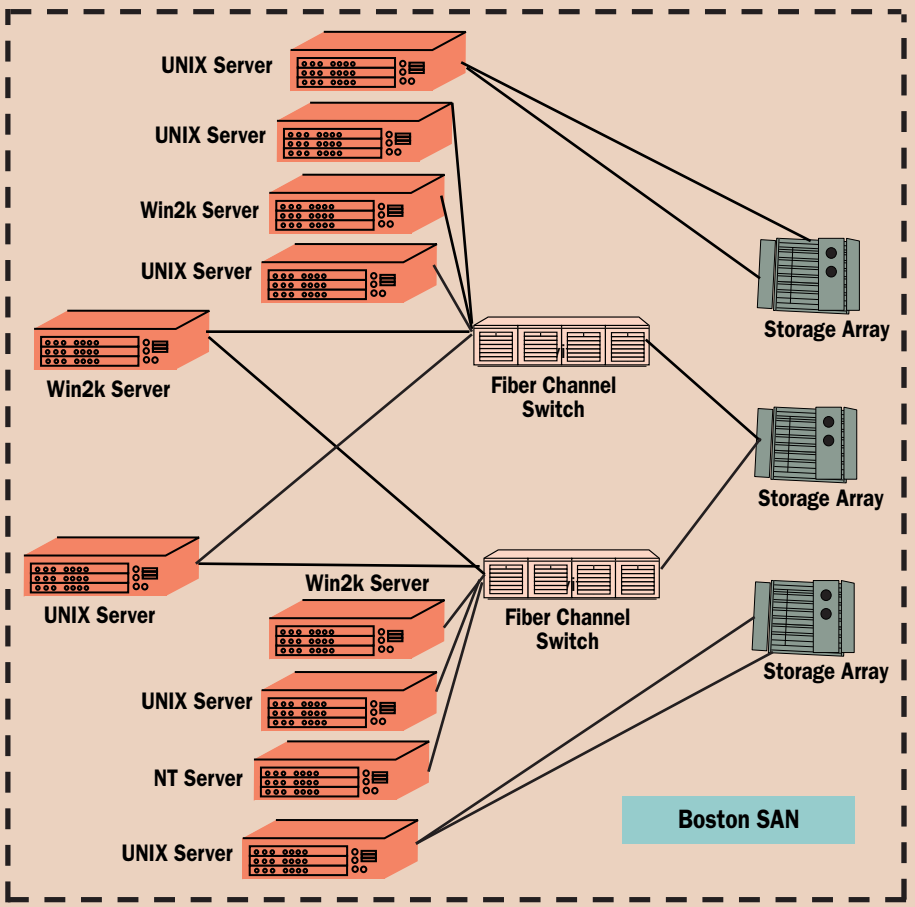

The NAS prices were nice, but our applications needed more privacy and flexibility

including storage vendors EMC, Hitachi, Network Appliance, Sun Microsystems and Winchester Systems. We also talked to SAN switch providers Brocade and Gadzoox. We knew we would be comparing, in some cases, an apple to an orange to a peach to a plum, but we believed each vendor had an interesting enough value proposition to merit consideration.

We wanted to implement a SAN and keep some direct-attached storage—but we didn't dismiss network-attached solutions without consideration. NAS is often half the price of a comparable SAN, and a NAS solution would have been compatible with our datacenter and corporate LAN. We also had good capacity planning, utilization and latency information, and knew that the existing infrastructure could have easily accepted the NAS traffic.

But we had performance and security worries about using NAS. The way that the NAS disk arrays are set up and configured makes them much more vulnerable to unauthorized access than a SAN, which is more private. Moreover, our team members who had previously worked with NAS expressed reservations about using it to support the applications that we were running.

They pointed out that, with NAS, we would not have the ability to differentiate between our own LAN traffic and the SAN traffic for the multiple customers we were hosting and their different types of data. In other words, NAS wouldn't let us meet the different risk requirements we had for each customer. Finally, we concluded that although NAS prices are nice, and NAS has made great strides to become competitive with SANs, NAS did not make sense for the types of applications that we were running.

We also considered staying 100 percent with DAS, and upgrading to Fibre Channel or ultra-SCSI. That would let us increase our I/O speed, but stay with the technology we knew was working and not rock the boat. But DAS wouldn't let us share storage between systems, so we wouldn't gain the efficiencies we were looking for, and we would have to continue physically handling the disk arrays and reconfiguring the volume management software for each server, each time we needed to make a storage change.

We quickly found ourselves looking at full Fibre Channel SAN architectures. Not only would a SAN be more private and perform better for our applications than the NAS, it would allow us to

Future concerns included disaster recovery and higher bandwidth

Getting Started With SANs

It seems obvious, but the first thing to do is figure out why you may need a SAN. Here are some of the questions that should be asked:

- Where is the data now that you are looking to move into the SAN?
- What are the main reasons that you feel you need a SAN?
- What is the size of the data volume you expect?
- When will the data arrive (scheduling)?
- What is your risk aversion level?
- Are there any special security concerns?
- How many employees do you have to spare to work on this project?
- What operational processes will you need to rethink before the SAN is installed?

These are detailed questions, but you really do need to understand exactly what data you have, where it is, how it will be moving and when. If you are running a Web cluster, email system, or customer relationship management (CRM) system, you'll find that you have different data storage and movement requirements□

limit our physical handling of the disks and even the volume management software changes in some cases.

To meet our security requirements, we kept a few direct-attached servers, which were Unix systems that we felt should not share the same physical fibers of the SAN or spindles in the disk drives with other data. Many SAN vendors will tell you that we could have achieved the separation we were looking for through "zoning" at the Fibre Channel switch level and through the storage management and volume management software. But we weren't confident that these mechanisms would provide sufficient protection. Moreover, any break-ins would have broken our service level agreements and possibly caused us other legal problems.

Anyone who subscribes to BugTraq and CERT can attest to the high level of attacks these days. Although we haven't seen or read about direct attacks on storage technology, this does not mean that it isn't taking place. As Lord Chesterfield said, "Judgment is not on all occasions required, but prudence is." When security is the question, prudence is the answer.

Once we had determined our design, as shown in Figure 1, we refined our requirements and put them out to a subset of the original group of vendors that we believed could provide a solution within the time frame we needed and with the technology questions satisfied. The non-disclosure agreements I have signed preclude me from further discussing the vendors by name. Ultimately, we selected two vendors as finalists, but our CEO influenced the final choice.

Refining The Design, Planning For The Future

As we finalized the SAN architecture, we wanted to maintain flexibility, both for near-term changes and for future growth. We addressed concerns about backups, disaster recovery and higher bandwidth between the servers and the storage, and felt our architecture gave us the future-proofing we were looking for. Our future state would have looked something like Figure 2.

We left our backup procedures on our datacenter LAN (not shown in the diagrams). From a capacity perspective, there was no pressing need to move to SAN backups, although the vendors assured us we could do so in the future. Many said we could simply attach our tape backup drives to the SAN, and then send the backups over the SAN instead of over the LAN. They told us this choice is popular with customers who are constrained on their datacenter LANs.

We figured if anything was going to bog down, it would be our initially-small SAN rather than our datacenter LAN, especially if backups from our many other servers hit the SAN at the same time as the heavy loads and updates our CRM systems required. Using a SAN for backups can also require additional software from some of the storage and volume management vendors.

Another future concern was the need to establish a disaster recovery (DR) site. We had always projected using our San Francisco datacenter as our DR site once the volume of business warranted this. In this regard, we talked with Cisco about using iSCSI router/gateways on our multihomed IP-VPN, and doing some mirroring of servers between our Boston and San Francisco sites. Besides the iSCSI routers, this would have required the same brand of storage systems at both ends, as well as vendor-specific DR backup software. Most of the storage vendors mentioned above either have proprietary DR software or partner with volume management vendors, such as Veritas, to offer this software, which can cost as much as a server or SAN switch (e.g., $50,000 to $80,000 depending on features and functionality).

Planning for higher bandwidth between the servers and the storage devices over the SAN fabric also was a future concern. We wanted to make sure that we were locking in our SAN for the next few years and that we would not have to upgrade too many major pieces, so we carefully considered our Fibre Channel switches and host bus adapters (HBAs), the storage networking equivalent of network interface cards (NICs).

Most SAN switch and storage device vendors have "qualified" HBA vendors they work with, and it makes sense to stick with these vendors. HBAs also interact with the volume management

FIGURE 2 Future SAN Architecture

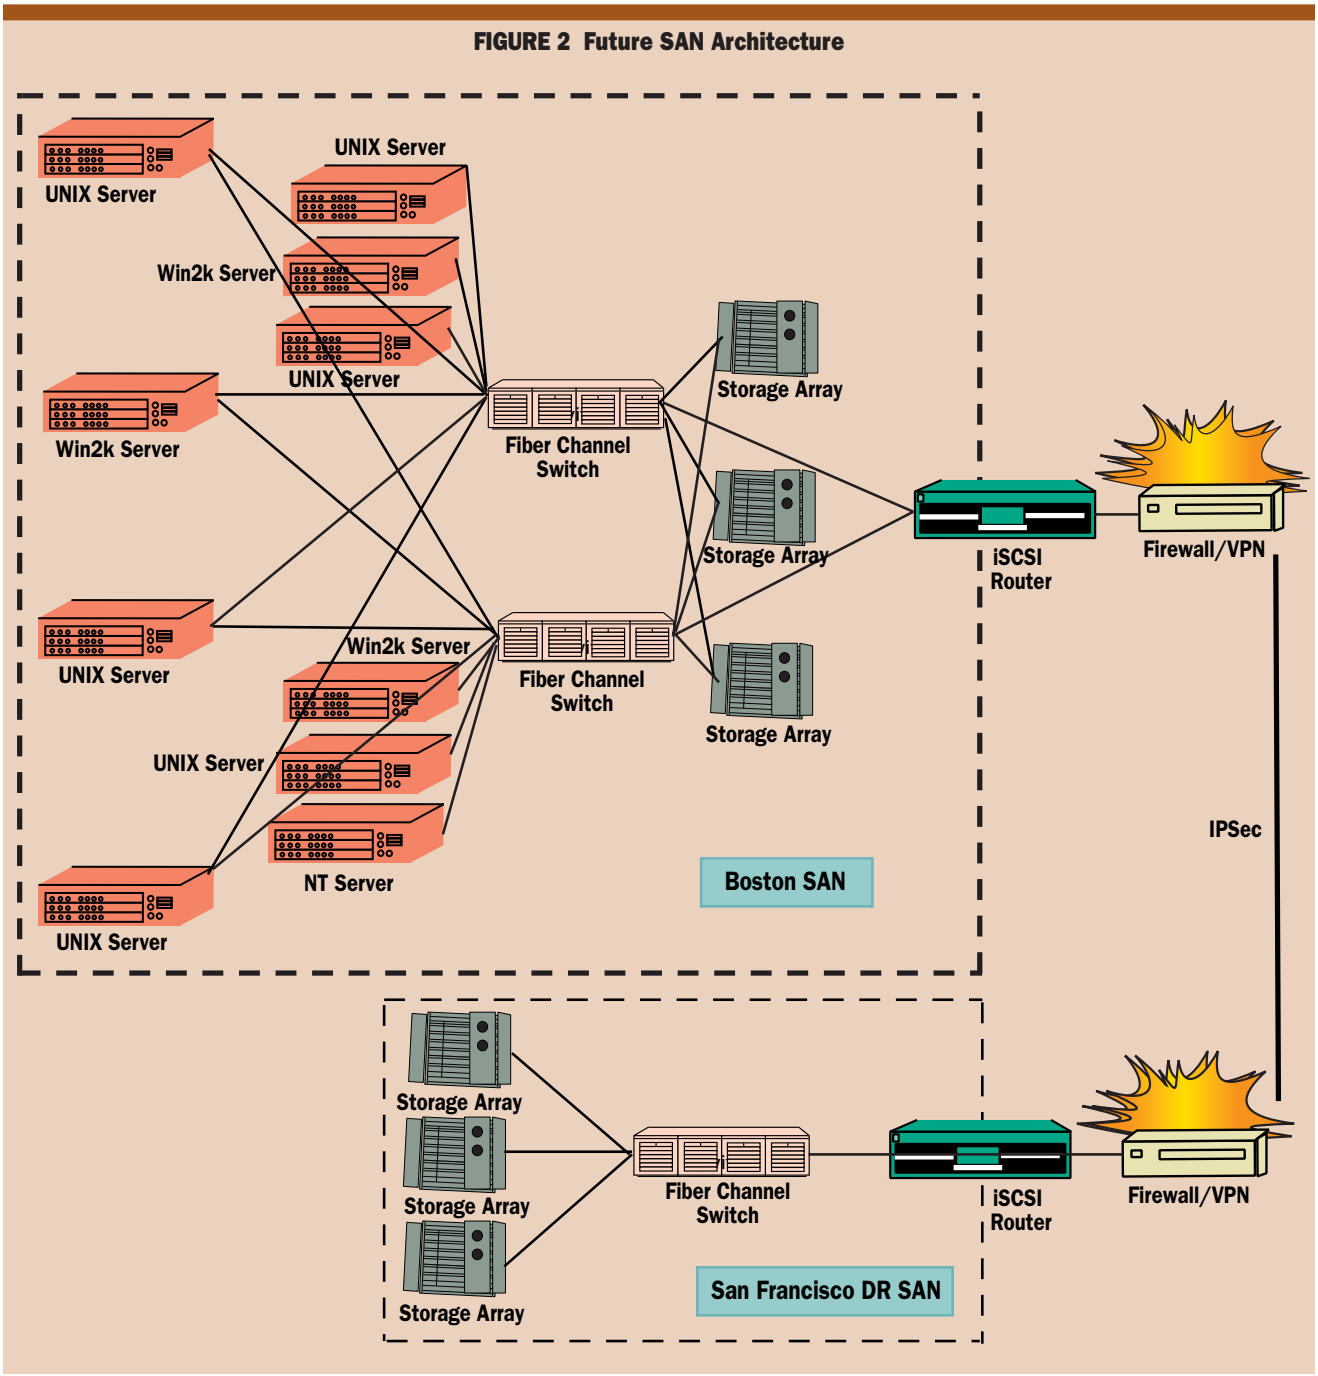

software, and you may need a version of it that is also HBA-qualified.

We also wanted to make sure that we were not locking ourselves into 1-Gbps fiber pipes, even though 2 Gbps were just being introduced. We were starting with 15, 1-Gbps fibers into the SAN switches and only 2–4, 1-Gbps fibers out of the switches to the storage arrays. We wanted our HBAs to be 2-Gbps capable, especially on the storage device side, and our SAN switches upgradeable to 2 Gbps. Going with the 2-Gbps HBAs in the beginning would mean that we would be set, from a SAN network fabric perspective. But this would also mean more expensive HBAs and an upgraded SAN switch that was 2-Gbps-ready. We ended up with 1-Gbps HBAs in the servers and a set of two, 1-Gbps-ready SAN switches that could be upgraded at a later date to 2 Gbps.

Sometimes our vendor representatives got confused when converting between storage and SAN capacity metrics. As with all networking technologies, SAN throughputs are quoted in bits per second, while storage, like servers and disk drives, is quoted in bytes. As a rule of thumb and to make the math easy, a 1-Gbps SAN switch is pushing about 100 Megabytes of data, once various types of overhead are taken into account. Only a few, so

SANs can help to increase the throughput of applications and storage devices

far, can push upwards of 140 Megabytes and use the 2-Gigabit HBAs.

This brings up another point: LAN and SAN switch technologies currently have much higher throughputs than the storage arrays. To compensate, some arrays buffer occasional overloads in memory, or they signal back to the servers to slow down. The arrays themselves are designed to trade off among port density, capacity and speed. Some have higher port density but less per-port storage, or they add larger, but slower disk drives.

Our SAN initially contained around 15 connected servers. We knew that once we brought the other two direct-attached servers and their storage devices into the SAN, we would need at least eight more ports available on our two SAN switches. This was the plan in the first year, assuming we could get our client contracts renegotiated such that the shared storage would meet their risk requirements.

From a network fabric standpoint we went with redundant switches, although we didn't opt for high availability (HA) mode. We had a mixture of dual-attached and single-attached servers, and, if one of the switches went down, we were willing to risk losing some of the servers.

Lessons Learned

We found that staffing and learning were bigger challenges than we had expected. Figure on about 1.5 times the number of people your storage vendor recommends, so that you can handle the inevitable hiccups in the installation process. Many of these are not anyone's fault *per se*, but details that are easily missed in the acquisition process. For example, you might have the wrong version of volume management software on one or more of your servers, causing a little downtime and the need for some different HBAs.

I have also seen and heard of more serious errors that have caused major downtime, such as misconfiguring storage systems, misunderstanding the storage layouts or forgetting to extend backup areas. With IT organizations being squeezed and employees doing more and more, sometimes steps in processes get missed.

You also can expect to need the vendor engineer onsite every so often over the first few months for knowledge transfer and to tune the system. All SANs run differently and need to be tuned and, in some cases, reconfigured, depending on the applications you are running. In the beginning the learning curve can be very steep and you will have to rely on your vendor for a lot of the heavy lifting.

Once we climbed the learning curve, we found the SAN was pretty easy to set up and to keep properly configured. We ran into some issues trying to determine whether the network, application or server organization should control the configuration. We decided it should be our server group, since the SAN storage technology went through the same change management process as the rest of our infrastructure, but we also crosstrained our networking group on the SAN switches and on configuration of the storage devices. This approach made both groups happy from a training and experience perspective, while not constraining anyone's daily activities.

Conclusion

SANs are not as complicated as the literature would have you believe, although vendors are still inclined to talk about the future and sell you the present. Make sure to understand your own requirements and stick to them.

If you are looking at multiple vendors, as you should, make sure that you come up with a common configuration for all the vendors to meet. You will never be comparing exactly apples to apples, but you want to avoid, say, justifying a watermelon compared to an apple or a plum. Even then—because IT expenses are being closely scrutinized and politics can enter the mix— the best technology does not always win.

SAN implementations are bound to pick up as the economy turns around and IT spending increases. They help companies increase the throughput of their basic applications and storage devices. Another boost in SAN spending will come from government regulations detailing how, where and for how long information needs to be accessible and archived.

From a market and technology viewpoint, consolidation is inevitable. We will have fewer vendors, and SAN technology will eventually be just another optional blade in a general-purpose LAN switch. Until then, however, SANs still have some maturing to do. They will probably grow more like Virtual LAN segments, with strict access control lists permitting traffic to traverse between servers and storage devices as well as between storage devices and other storage devices, and more products will be developed that scale from small to medium to large□

Companies Mentioned In This Article

Brocade Communications (www.brocade.com)
BugTraq (www.securityfocus.com)
CERT (www.cert.org)
Cisco (www.cisco.com)
EMC (www.emc.com)
Gadzoox Networks (www.gadzoox.com)
Hitachi America (www.hitachi.com)
Network Appliance (www.netapp.com)
Sun Microsystems (www.sun.com)
Veritas Software (www.veritas.com)
Winchester Systems (www.winsys.com)