Next Generation Datacenter SANs

An IDC White Paper

Analysts: John McArthur, Robert Gray, Eric Sheppard, and Bill North

EXECUTIVE SUMMARY

Datacenters with multiple SAN implementations today will find two advances in SAN technology important and useful. These will enable next-generation datacenter SANs to further consolidate storage and improve manageability.

IDC believes that consolidation can take can take place in the storage network fabric as well as in storage arrays. For example, multi-level multi-protocol switches with higher port counts can consolidate the switching function within the SAN, as well as offer end users an alternative to separate, purpose built bridging devices. Consolidation may also occur through the use of larger switches, i.e., switches with increase port density. Using larger building blocks within the SAN will result in fewer switches and also reduce the number of ports needed for inter-switch communications.

Manageability will improve due to advances in storage and network virtualization, the technique of mapping physical characteristics of devices to logical attributes that are simpler for operators to manage. IDC believes that standards play an important role in advancing such virtualization within the multi-vendor world of enterprise SANs. Co-operative development efforts at the Storage Network Industry Association (SNIA) are expected to ultimately generate an industry standard vocabulary that becomes the basis for cross-vendor virtualization efforts.

Longer-term IDC believes that the concept of utility computing — the idea that information systems can be provisioned like electrical and telecommunication utilities — will mature within the data center, and evolve to include more applications that aren't necessarily considered datacenter applications today. Increased consolidation and improved manageability of storage systems are essential steps in the direction of increasing the reach of utility computing. Along with other steps (e.g., the development of blade servers designed for consolidated processing that can be provisioned), the IT industry in general is heading in the direction of the utility model — albeit at a conservative pace.

Sponsored by Cisco

IDC
Analyze the Future

5 Speen Street • Framingham, MA 01701 USA • Phone 508.872.8200 • Fax 508.935.4015 www.idc.com

INTRODUCTION

Important changes are afoot in datacenter storage technologies. Today's storage area networks (SANs) have succeeded in whetting the appetite of IT consumers who want to consolidate storage to improve utilization and provision shared storage resources. Tomorrow's datacenter storage networks will provide greater consolidation, will be easier to manage, and will expand to include more applications that are not traditionally found within the datacenter. Technology improvements underlie these trends.

Today's Datacenter SAN

Today's datacenter SAN consolidates and allocates storage resources to a collection of servers that are typically running demanding applications. SANs are also deployed as shared infrastructure in organizational units such as divisions, departments or campuses. Within the datacenter, SAN adopters are taking the first steps toward a new model of storage — a utility model where storage, like electrical power, network services and cooling, is provisioned. IDC notes, however, that many large data-centers often have several SANs, each dedicated to supporting different application or organizational needs. Within these environments, SANs are not integrated with one another or with workgroup and mid-size systems in the enterprise.

Today's Datacenter SAN: Benefits & Challenges

SAN technology provides a viable approach to gaining control over large amounts of storage previously attached directly to servers. Because storage managers can reallocate shared resources, greater utilization becomes a significant benefit. The pool of shared storage can be a heterogeneous mix of tape and disk arrays, which can be provisioned as required for different enterprise applications attached to the SAN. Backup and recovery routines can be woven into the SAN.

In many large data-centers, IDC believes that multiple SANs are maintained separately for different reasons (see call out section entitled When SAN Integration Adds Value). Today's switches and directors allow SANs to scale to serve hundreds of servers and storage devices. Though SANs can theoretically scale to thousands of nodes, few ever do. Though data-centers with a very large number of nodes certainly exist, they often contain multiple, separate SANs.

Ease of use is a considerable challenge for SAN managers today. In fact, IDC research indicates that ease of use is the most prevalent concern for owners of SANs, and that this concern grows as the size, scale and complexity of the SAN increases. This is due, in part, to the fact that today's SANs use fairly new technologies that often

do not leverage industry standards. Further, today's SANs require retraining of IT staff that typically have little experience with network storage protocols. Also, a detailed understanding of the SAN elements — the switches, host-bus adapters, and various storage devices — is needed in order to maintain the SAN.

SAN TECHNOLOGY IN A NUTSHELL

A storage area network (SAN) consists of a dedicated network that interconnects servers and their applications to storage resources with the use of a switch or hub. Today, market preference for hubs has largely gone away, as many switch prices have declined to a price per port that was once reached only by low-cost hubs with limited functionality.

Host-bus adapters installed in servers look to the SAN for blocks of data as if it were directly attached. Bridges are used to provide gateway functionality, extending SANs over disparate protocols outside the dedicated network. Storage devices are ordinarily a mix of disk and tape resources with different performance characteristics. The switch routes the data between servers and storage devices. Often these switches are relied upon to ensure the availability of the fabric within a SAN.

Early adopters of SANs were large data centers with commensurately large storage management needs and budgets. As such, risks at this level were often quite high. SANs are also employed by mid-size data centers or workgroups with modest storage needs and budgets. These SANs have tended to be less complex, often supporting less critical applications.

Storage system managers who oversee SANs are among the most skilled in the IT industry, and they need to be. SANs are complex and a deeper understanding of both network and storage technologies is needed to operate them. For managers who come from the mainframe world, SANs are a return to the good old days of centralized control. For managers who come from the distributed systems world, consolidation in general, and SANs in particular, are approached with caution. In fact those with a background in distributed environments often limit the exposure of a SAN by deploying SANs on a per project or per application basis.

Tomorrow's Datacenter SANs

Next-generation SAN technologies will provide data-centers with greater scalability to enable further consolidation of storage resources and improved ease of use.

Further Consolidation

Large data centers containing multiple SANs will continue to consolidate storage by improving the inter-SAN communication within the datacenter. This opens the possibility of greater storage consolidation and optimization than can be achieved in individual, small-scale

storage networks. In addition, large datacenter SANs will become more accessible to mid-size applications that today are running on smaller servers with attached storage in departments and workgroups.

Ease of Use

In the networked storage industry, efforts to make SANs easier to use fall under the virtualization umbrella. Virtualization is the use of software to mask low-level system details so that SANs are easier to manage. Moreover, virtualization provides a platform for additional software that automates typical SAN activities so that SANs increasingly manage themselves.

SAN TECHNOLOGY TRENDS

Consolidation and manageability improvements are two major trends that IDC expects to have a significant impact on datacenter SANs. These installations, with terabytes of storage under management, are the early adopters of the consolidated network storage model. As such, the needs that this community has encountered while deploying, scaling, and managing SANs are harbingers of trends to come throughout the storage industry.

Continued consolidation

From the onset, SANs have provided an architecture for consolidating storage devices. Disk arrays with different performance and capacity characteristics can be integrated with tape libraries to provide a single pool. Ultimately SANs have resulted in simplifying management by provisioning from a one pool of storage resources.

Though one major driver of SAN deployment within a datacenter has been consolidation, the need to prevent the SAN from becoming unmanageably complex, as well as the desire to avoid inefficient use of ports as inter-switch links has limited the degree to which many datacenter SANs have consolidated. IDC believes that such datacenter SAN deployments will eventually consolidate further by leveraging switches and directors with larger port counts. These larger SANs will be assembled with fewer switches and each switch will manage a larger number of ports. Ultimately these types of SANs will have fewer components to manage.

High-performance switches are another catalyst for consolidation. Be it at 2-Gb/s wire speeds or by trunking multiple physical ports to create a single, high-speed, logical connection, faster Fibre Channel ports can be used to join SANs in a datacenter-wide fabric. This allows operators to gain views of the entire data-center's operations from a single console.

The Coming of Multi-Protocol Switches

Ample opportunity exists to reduce the number of devices needed to support multiple protocols within a single SAN.

Today, the use of multiple protocols within datacenter SANs is not uncommon. Frequently, these SANs encapsulate the Fibre Channel protocol into IP packets, then send the packets long distances over secure, high-speed networks. Such functionality is provided by a purpose built device often not built by a switch supplier. IDC believes that incorporating or migrating encapsulation function to the SAN switch and making it a portion of a single supplier solution would not only reduce the number of devices being managed, but also reduce the complexity associated with having to rely on disparate suppliers for switching and encapsulation functions.

IDC notes that switches with incorporated encapsulation functionality will compete with traditional purpose built devices. For those that choose to leverage purpose built devices, benefits of further consolidation are by no means out of reach. In fact for these environments, users can utilize a central switch as an efficient way of locating gateways between protocols, since one gateway can be used as the central point rather than locating gateways at multiple smaller switches or at even higher numbers of servers.

The need for encapsulating one protocol over another, such as Fibre Channel encapsulated in IP, may subside over time as the introduction of competing SAN protocols (e.g., iSCSI) allow for long distance data transmission natively. That said, this will happen over time, after such protocols mature and software is rewritten to accommodate them.

It is IDC's belief that users will continue to use bridging devices to interconnect SANs. That said, the market will soon have the added choice of a multi-protocol switch with some portion of ports accommodating long-reaching protocols such as ethernet at Layer 2, and IP at Layer 3.

Lastly, IDC believes that it is important to note that the end-point for consolidation is an eventual move to a utility model for storage, networking and processing — that is, utility computing or "IT on demand." We believe that the utility model is viable, but that progress toward that model will be incremental and last over many years to come.

More on Multi-level Multi-protocol Support

IDC expects that enterprises will continue to utilize LANs and the WAN backbones with different protocols at different layers, particularly layers 2 and 3 in the ISO Reference Model. As SAN resources reach out more broadly in the enterprise, SANs will need to support these different protocols and levels as well.

- At level 3, the protocol of choice is the Internet Protocol (IP). Competitors to IP do exist, but are of vanishing importance.
- At level 2, Ethernet, Asynchronous Transfer Mode (ATM) and Frame Relay are the common protocols that co-exist in intra-enterprise and extra-enterprise networks.

It is therefore logical to assume that a multi-protocol SAN switch would have the various protocols of the SAN world, of which Fibre Channel is currently the largest, and also those of the network world,

of which ethernet and IP are largest. A SAN switch attached to the LAN can become a shared storage resource for the enterprise. Extending this argument, a SAN switch that can be attached to metropolitan or more broadly all wide-area networks can potentially be used by many sites.

Extended connectivity could be particularly attractive in MANs, where many service providers are considering all-ethernet configurations, riding over fiber. Thus SAN switches that use ethernet are attractive both in the LAN and the MAN. More complex networks would require use of IP as well. Broadening this further, the use of ATM and Frame Relay would allow the use of SANs with other popular networks. Today, bridges and edge devices provide linkages to different network protocols. In the future, multi-protocol SAN switches will evolve to compete in satisfying this need.

Improving SAN manageability

IDC believes there to be three basic components that can help to improve SAN management: virtualization, increased standardization, and increased partitioning technologies.

Storage virtualization

IDC defines storage virtualization as the use of software to simplify the management of complex storage configurations by using logical representations of physical resources.

For example, virtualization allows a storage manager to request manually that data supporting a critical enterprise application be backed up to tape. Typically the operator executes this request from a workstation that provides a high-level view of the SAN environment. The operator need not know the exact location of the dataset or the availability of a tape drive and a scratch tape. Such details are held in a repository and are called upon as needed when the operator issues a high-level command.

With the addition of a rule or policy engine, the process of backing up application data can be automated. The operator sets into place the conditions for backup and the storage system automatically follows those instructions, provides status reports and alerts as needed.

IDC believes that, over time, virtualization will make networked storage products easier to manage and provide the foundation for automation products that require less management attention than their manual alternatives. In a nutshell, storage managers will be asked simpler questions and fewer of them.

Increased standardization

The adoption of standards is a sign of maturity in an industry. The Storage Networking Industry Association (SNIA) is one of the leading forums for cooperative development work. SNIA's Common Information Model (CIM) and the Web Based Enterprise Management (WBEM) technologies are emerging as the industry's defacto standards. The CIM standard is a vocabulary for describing storage

devices and the processes related to delivering storage services in a heterogeneous network environment. WBEM provides a standard network architecture for storage management.

More recently CIM and WBEM have been used as the basis for a new management interface called Bluefin. With support from 17 member companies, SNIA is well positioned to assist the industry in focusing on a single management interface. IDC believes that SAN manageability will improve as more companies adopt uniform access methods and protocols for managing components that make up the SAN.

Improved network partitioning technologies

For enterprise SANs to become more widely accessible in an enterprise, robust technologies must be available to partition the SAN for security and scalability purposes. With regard to security, these partitioning techniques will need to stand up to formal audits when the security of sensitive information is at stake. Authorization and authentication technologies and policies beyond the SAN will need to meet these same stringent tests. Users will expect parts of the SAN to behave as if they are entirely isolated from one another, that is, as if the data were in its own dedicated storage system.

Partitioning is important to SAN scalability as well. Topology changes, for example adding a switch or reconfiguring pooled storage resources, trigger messages among SAN devices and, for a time, SAN services are unavailable. Partitioning allows operators to limit the scope of this message and discovery process and avoids potential disruption and downtime in parts of the network that are not directly affected by the changes. This ultimately reduces the amount of management intervention and increases overall storage network and system availability.

The raw ingredients for partitioning exist in different SAN components. Zoning capabilities provide a low-level foundation for partitioning and performance tuning in the future. IDC expects to see virtualization techniques applied to the partitioning challenge and the result will be ways to subdivide a single, large physical SAN into multiple logical SANs, each provisioned to meet a particular set of application needs. Such virtual SANs will operate independently and provide the partitioning necessary as enterprise access increases.

WHEN SAN INTEGRATION ADDS VALUE

Types of enterprise data are not created equal.

Following the hierarchy of most organizations, data may be characterized as enterprise level, line-of-business and departmental, and personal. These types of data are quite different.

- Enterprise data is information that has value throughout the organization. A good example is shared access to product support and customer care data so that engineering can design better products, marketing can establish new needs, and manufacturing can monitor quality.

- Throughout the enterprise, datasets are cached in support of Web services, for example, or in print queues and replicating email servers.
- Some departments (e.g., human resources, legal, and finance) house data that is sensitive. Access to such data is generally limited to "need-to-know" — even within the department.
- New requirements highlight the importance of maintaining privacy for personal data concerning both employees and customers

The advantages of wide access to a SAN correlate dramatically in accordance with data type.

- Data of widespread interest to the corporation is valuable to share across traditional organizational lines. In addition to customer care, there is commonly widespread interest in product and service definitions, corporate communications to customers as well as employees.
- Cached and replicated data need not be widely accessible from a common store since this data is neither unique nor long-lived.
- Departmental data is, by definition, primarily of interest within each department. Provisioning departmental storage from an enterprise SAN will meet with significant resistance since IT systems have traditionally mirrored the organizational hierarchies in which they are embedded.
- Personal data and sensitive data must be protected in ways that may be audited or required to meet legal requirements. While secure partitioning of SANs may be technically defensible, it remains simpler in many cases to keep such data in an isolated network.

Major market hurdles/issues

Throughout much of the IT world, SANs have a reputation for being complex, expensive and difficult to manage. As a result, the technology is deployed only when necessary and many times with reliance on a third-party expert. In addition to delivering quality products that are simpler to manage, suppliers of networked storage will need to convince the skeptics that these systems really are easier to install and use. Further, to be attractive to a large portion of the total available market, SANs must, over time, mature to the point that third-party experts are relied upon for only the most complex configurations.

IDC believes that similar challenges will arise when SANs are more widely integrated into enterprise networks. Resource sharing, for example, is a difficult benefit for organizations to quantify and factor into an ROI or TCO calculation.

IDC ANALYSIS: OPPORTUNITIES AND CHALLENGES

Opportunities

Users of datacenter SANs have realized the value of consolidated storage and, over time, will develop an interest in further integration of SANs with other network services in the enterprise.

Standardization and automation (virtualization) offer great promise both for existing SAN customers and especially for next-generation SAN customers who are waiting for systems that are easier to deploy and maintain and for which standards make investments safer.

Improvements in the throughput and port capacity of the SAN switch can and will accelerate the further proliferation of SAN technology for large enterprises with a high demand for shared data.

Challenges

The transition from non-standard to standard components is difficult for any industry, and the storage industry is no exception. Traditionally, first-movers are punished when they release products with plug-compatible options. Regardless, it must occur.

The virtualization movement is clearly aimed in the right direction — to make networked storage simpler to configure, deploy and manage. Setting goals and achieving them are two different events, however, and the virtualization effort will be measured by customer response to new systems that aim to be both flexible and simple to operate.

IDC believes that there are a considerable number of large data-centers that have multiple SAN islands. While we believe there to be opportunity for further consolidation within many of these environments, we again note that not all data is equal and that a island approach may sometimes be the best.

Lastly, we note that larger, faster, multi-protocol SAN switches must succeed in simplifying the SAN environment or face slow adoption.

CONCLUSION

Progress will continue unimpeded in the networked storage industry. IDC encourages consumers, particularly those who operate multiple SANs in their data centers, to monitor early adopter behavior as suppliers offer technology that enables further consolidation of storage resources and advancement in manageability.

IDC Worldwide Offices

CORPORATE HEADQUARTERS

IDC
5 Speen Street
Framingham, MA 01701
United States
508.872.8200

NORTH AMERICA

IDC Canada
36 Toronto Street, Suite 950
Toronto, Ontario M5C 2C5 Canada
416.369.0033

IDC California (Irvine)
18831 Von Karmen Avenue
Suite 200
Irvine, CA 92612
949.250.1960

IDC California (Mountain View)
2131 Landings Drive
Mountain View, CA 94043
650.691.0500

IDC New Jersey
75 Broad Street, 2nd Floor
Red Bank, NJ 07701
732.842.0791

IDC New York
2 Park Avenue
Suite 1505
New York, NY 10016
212.726.0900

IDC Texas
100 Congress Avenue
Suite 2000
Austin, TX 78701
512.469.6333

IDC Virginia
8304 Professional Hill Drive
Fairfax, VA 22031
703.280.5161

EUROPE

IDC Austria
c/o Loisel, Spiel, Zach Consulting
Mayerhofgasse 6
Vienna A-1040, Austria
43.1.50.50.900

IDC Benelux (Belgium)
Boulevard Saint Michel 47
1040 Brussels, Belgium
32.2.737.76.02

IDC Denmark
Omøgade 8
Postbox 2609
2100 Copenhagen, Denmark
45.39.16.2222

IDC Finland
Jarrumiehenkatu2
FIN- 00520 Helsinki
Finland
358.9.8770.466

IDC France
Immeuble La Fayette 2
Place des Vosges Cedex 65
92051 Paris la Defense 5, France
33.1.49.04.8000

IDC Germany
Nibelungenplatz 3, 11th Floor
60318 Frankfurt, Germany
49.69.90.50.20

IDC Italy
Viale Monza, 14
20127 Milan, Italy
39.02.28457.1

IDC Netherlands
A. Fokkerweg 1
Amsterdam1059 CM, Netherlands
31.20.6692.721

IDC Portugal
c/o Ponto de Convergancia SA
Av. Antonio Serpa 36 - 9th Floor
1050-027 Lisbon, Portugal
351.21.796.5487

IDC Spain
Fortuny 18, Planta 5
28010 — Madrid
Spain
34.91.787.2150

IDC Sweden
Box 1096
Kistagangen 21
S-164 25 Kista, Sweden
46.8.751.0415

IDC U.K.
British Standards House
389 Chiswick High Road
London W4 4AE United Kingdom
44.208.987.7100

LATIN AMERICA

IDC Latin America
Regional Headquarters
8200 NW 41 Street, Suite 300
Miami, FL 33166
305.267.2616

IDC Argentina
Trends Consulting
Rivadavia 413, Piso 4, Oficina 6
C1002AAC, Buenos Aires, Argentina
54.11.4343.8899

IDC Brazil
Alameda Ribeirao Preto, 130
Conjunto 41
Sao Paulo, SP CEP: 01331-000 Brazil
55.11. 3371.0000

International Data Corp. Chile
Luis Thayer Ojeda 166 Piso 13
Providencia
Santiago, 9, Chile
56.2.334.1826

IDC Colombia
Carerra 40 105A-12
Bogota, Colombia
571.533.2326

IDC Mexico
Select-IDC
Av. Nuevo Leon No. 54 Desp. 501
Col. Hipodromo Condesa
C.P. 06100, Mexico
525.256.1426

IDC Venezuela
Calle Guaicaipuro
Torre Alianza, 6 Piso, 6D
El Rosal
Caracas, Venezuela
58.2.951.1109

CENTRAL AND EASTERN EUROPE

IDC CEMA
Central and Eastern
European Headquarters
Male Namesti 13
110 00 Praha 1
Czech Republic
420.2.2142.3140

IDC Croatia
Srednjaci 8
1000 Zagreb
Croatia
385.1.3040050

IDC Hungary
Nador utca 23
5th Floor
H-1051 Budapest. Hungary
36.1.473.2370

IDC Poland
Czapli 31A
02-781 Warszawa, Poland
48.22.7540518

IDC Russia
Suites 341-342
Orlikov Pereulok 5
Moscow, Russia 107996
7.095.975.0042

MIDDLE EAST AND AFRICA

IDC Middle East
1001 Al Ettihad Building
Port Saeed
P.O. Box 41856
Dubai, United Arab Emirates
971.4.295.2668

IDC Israel
4 Gershon Street
Tel Aviv 67017, Israel
972.3.561.1660

IDC South Africa
Building 9, Pebble Beach
Fourways Golf Park
Roos Street
Fourways, Gauteng
South Africa
27.11.540.8000

IDC Turkey
Tevfik Erdonmez Sok. 2/1 Gul
Apt. Kat 9D
46 Esentepe 80280
Istanbul, Turkey
90.212.275.0995

ASIA/PACIFIC

IDC Singapore
Asia/Pacific Headquarters
80 Anson Road
#38-00 IBM Towers
Singapore 079907
65.6226.0330

IDC Australia
Level 3, 157 Walker Street
North Sydney, NSW 2060
Australia
61.2.9922.5300

IDC China
Room 611, Beijing Times Square
88 West Chang'an Avenue
Beijing 100031
People's Republic of China
86.10.8391.3610

IDC Hong Kong
12/F, St. John's Building
33 Garden Road
Central, Hong Kong
852.2530.3831

IDC India Limited
Cyber House
B-35, Sector 32, Institutional
Gurgaon 122002
Haryana India
91.124.6381673

IDC Indonesia
Suite 40, 17th Floor
Jakarta Stock Exchange
Tower 2, Jl. Jend. Sudirman Kav. 52-53
Jakarta 12190
6.221.515.7676

IDC Market Research (M) Sdn Bhd
Jakarta Stock Exchange Tower II
17th Floor
Jl. Jend. Sudirman Kav. 52-53
Jakarta 12190
62.21.515.7676

IDC Japan
The Itoyama Tower 10F
3-7-18 Mita, Minato-ku
Tokyo 108-0073, Japan
81.3.5440.3400

IDC Korea Ltd.
Suite 704, Korea Trade Center
159-1, Samsung-Dong
Kangnam-Ku, Seoul, Korea, 135-729
822.551.4380

IDC Market Research (M) Sdn Bhd
Suite 13-03, Level 13
Menara HLA
3, Jalan Kia Peng
50450 Kuala Lumpur, Malaysia
60.3.2163.3715

IDC New Zealand
Level 7, 246 Queen Street
Auckland, New Zealand
64.9.309.8252

IDC Philippines
703-705 SEDCCO I Bldg.
120 Rada cor. Legaspi Streets
Legaspi Village, Makati City
Philippines 1200
632. 867.2288

IDC Taiwan Ltd.
10F, 31 Jen-Ai Road, Sec. 4
Taipei 106
Taiwan, R.O.C.
886.2.2731.7288

IDC Thailand
27 AR building
Soi Charoen Nakorn 14,
Charoen Nakorn Rd., Klongtonsai
Klongsan, Bangkok 10600
Thailand
66.02.439.4591.2

IDC Vietnam
Saigon Trade Centre
37 Ton Duc Thang Street
Unit 1606, District-1
Hochiminh City, Vietnam
84.8.910.1233; 5

IDC is the foremost global market intelligence and advisory firm helping clients gain insight into technology and ebusiness trends to develop sound business strategies. Using a combination of rigorous primary research, in-depth analysis, and client interaction, IDC forecasts worldwide markets and trends to deliver dependable service and client advice. More than 700 analysts in 43 countries provide global research with local content. IDC's customers comprise the world's leading IT suppliers, IT organizations, ebusiness companies and the financial community. Additional information can be found at www.idc.com.

IDC is a division of IDG, the world's leading IT media, research and exposition company.

02-315HARDWA3455
August 2002

IDC
Analyze the Future

www.idc.com