DOCUMENT RESUME

ED 417 707 IR 018 774

TITLE	Network Technology Based Application.
INSTITUTION	Northwest Regional Educational Lab., Portland, OR.
PUB DATE	1998-00-00
NOTE	19p.; Session 5 of "Information Technology in Education and Training (IT@EDU98). Proceedings"; see IR 018 769. For sessions 1-6, see IR 018 770-775.
PUB TYPE	Reports - Descriptive (141) -- Speeches/Meeting Papers (150)
EDRS PRICE	MF01/PC01 Plus Postage.
DESCRIPTORS	*Computer Networks; Databases; *Distance Education; *Educational Technology; Foreign Countries; Higher Education; Information Technology; *Telecommunications
IDENTIFIERS	Network Management; Network Protocols; University of Hawaii; Vietnam

ABSTRACT

The fifth session of IT@EDU98 consisted of four papers and was chaired by Nguyen Thanh Son (University of Technology, Ho Chi Minh City, Vietnam). "Distance Education at University of Hawaii" (David Lassner) investigates the technologies in use at the University of Hawaii, including: SkyBridge; HITS (Hawaii Interactive Television System); I-Net (Institutional Network); compressed digital video; and satellite. "An Approach to Distance Education by Using Network Technology" (Dam Quang Hong Hai) presents one method for creating a distance course by using network technology that can be implemented with the telecommunications network in Vietnam. "About the Ways To Solve the Shortage of IP Addresses" (Phan Cong Vinh) suggests ways for network managers to handle the high demand and short supply of IP addresses. "Introduction to Very Large Databases" (Do Hoang Cuong) defines a very large database, looks at maintenance issues, and provides guidelines for implementing solutions. (SWC)

SESSION 5

Friday, 16 January 1998

Session 5: Network technology based application

Chair:
Dr. Nguyen Thanh Son, University of Technology, HCMC, Vietnam

5-1. Distance Education in University of Hawaii
Dr. David Lassner, University of Hawaii, USA

5-2. An Approach to Distance Education by using network technology
Dam Quang Hong Hai, University of Natural Sciences, HCMC, Vietnam

5-3. About the ways to solve shortage of IP address
Phan Cong Vinh, Vietnam Post and telecoms Institute of Technology

5-4. Introdution to a Very Large Database
Do Hoang Cuong, University of Natural Sciences, HCMC, Vietnam

DISTANCE EDUCATION SYSTEM AT UNIVERSITY OF HAWAII

David Lassner, Hae Okimoto
University of Hawaii, USA

The term " distance education " and HITS (Hawaii Interactive Television System) have been floating round our campuses for several years. Some people may still be a bit puzzled as to what they really mean or if it should matter to them. Others may even have taught a distance education class using HITS and/or cable television without being aware of what else is possible. Hopefully this article will help clear up the airways a bit.

The term " distance education" usually indicates instruction which occurs when students are physically separated from their instructor. This is most often accomplished with the aid of telecommunications technology. Typically, the instructor is situated in one location, while the students may be in multiple locations (including in their home) on different islands. Distance education has been primary use of the University of Hawaii (UH) video networking systems. However with budget cuts and the ease and effectiveness of telecommunications technology, there have been increased requests for use of the systems for meeting and even accreditation visits. So, what are these video systems? The following are brief descriptions of the technologies currently in use.

SkyBridge- SkyBridge is Maui Community College's (Maui CC) microwave system which serves the three islands of Maui county. SkyBridge provides one channel of 2-way video among Maui CC and its education canters on Moloka'i, Lana'i, and in Hana. This provides residents in these locations with access to Maui CC and other courses taught on that campus. Skybridge was built with support from the Federal government.

HTTS- The Hawaii Interactive Television, uses microwave and IFTF(Instructional Television Fixed Service) transmission technology to provide 4 channels of video and audio communication among the islands. HTTS programs may utilize both 2-way video or 1-way video with return audio only. Since mid-1990, the largest use of HITS has been for the delivery of credit programs between the UH campus. The Hawaii Department of Education (DOE) also makes extensive use of HITS for direct instruction and teacher in-service training programs. In January 1995, stewardship of HTTS was transferred from Hawaii Public Television of Hawaii to the University of Hawaii.

Cable TV- As part of their cable franchise agreement, all commercial cable companies within Hawaii provide access channels which may be used for education programming. Most cable companies can receive live programming via HTTS, thus providing UH and DOE with nearly statewide live cable programming capabilities.

Programs on cable are either reproduced to be shown as a scheduled time or broadcast live with return audio capabilities via telephone.

I-Net(Institutional Network)- Cable franchise agreements mandate that cable companies help the State develop as internal infrastructure by providing fiber optic cables and/or other telecommunication services to specific State locations, usually at cost. Fiber connecting O'ahu campuses allows another option for delivering video between campuses. For example, Leeward Community College (Leward CC) cannot originate live video programming directly on the HTTS system. However, since they are a video I-Net to UH Manoa, where the programs can be switched onto HTTS. The I-Net video system transmit broadcast-quality video, using the same technology Oceanic Cable uses to distribute video around the island.

Compressed Digital Video- UH is now beginning to use compressed digital video, often referred to as videoconferencing technology. This differs from cable TV and the original HITS system in that the video system is digitized and compressed before being transmitted. This permits the transmission of many more "channels" the original analog HITS and SkyBridge technologies. There are two compressed digital video systems currently in use by UH. The VideoConnect pilot project connects eight UH and DOE sites on six islands in order to test this proposed new GTE Hawaii with Tel service. VideoConnect now serves as the primary means to connect UH WestHawai'i with UHHilo and Hawaii Community College. We have also installed videoconferencing equipment (from Compression Labs Inc.,or CLI) at several UH O'ahu sites. This system has been used to provide Leeward CC classes at our Leeward CC at Wai'anae education center. This is the same technology installed by the State of Hawaii for their statewide videoconferencing service.

Satellite - While the systems already described carry video signals within the State, satellite technology is the most common means for receiving programs from outside Hawaii. Satellite teleconferences are generally received on a UH satellite dish and carried using one or more of the above systems to one or more locations where they can be viewed live. Audio interaction with the presenters is usually possible by calling a toll-free telephone number. Many campuses have satellite down-link facilities; most programs are received either on a campus dish or on the down-link facilities of the UH Learning Center (LTRLC) or at Hawaii to the Mainland and Asia/Pacific regions.

During the fall we hope to use additional technologies for video transmission. We will be installing desktop video systems in distance education facilities on different islands to permit more cost-effective small group and individual interaction, such as electronic office hours. We are also testing dialup videoconferencing, as a less expensive alternative to satellite for out-of-state video connection.

In order to improve the UH modems, we are adding 100 additional lines, which will take our total over 250 lines. These will include higher speed modems and an improved management capability which will reduce our staff workload and provide more options for managing this increasingly scarce resource. We are planning to partition some modems into "express" modem pool to make it easier for people to be able to get in just to check email quickly. Because of the financial difficulties, we will not be able to keep up with the demand for free dialup services. The UH community consists of some 60,000 faculty, staff, and credit students. To serve this population with high quality service, it would require at least a 10-fold expansion of our modem pool, with associated capital and recurring costs. We are pursing several ways to reduce the cost of supporting dialup modems. As an example, at the current market price of about $25/month for unlimited dialup access it would cost $18 million a year for 6000 users. Obviously that 's not the real cost, but it conveys a sense of the scale of the problem we face.

Fortunately, there are now many private ISP who manages their own dialup modem pools and sell services to the public. We are now implementing direct connections with local providers through a project we call the Hawaii Internet Exchange, or HIX. This will provide improved service for any member of the UH community who chooses to buy service from a private provider but still want to reach UH resources. It will also improve connections from UH to campuses to information services hosted by the private provider.

Many universities are giving up and outsourcing dialup access by allowing a private provider to sell access directly to students and faculty. And some are beginning to charge to recover costs. We could ration access by limiting cumulative usage, perhaps to 20 hours per user per month, or we could begin to charge faculty, staff, and students for dialup access. None of these solutions are very appealing, but it is clear to anybody who thinks about it that either funding patterns or expectations have to change. The basic problem is that dialup costs are roughly linear with usage, dialup usage is growing exponentially, and institutional budgets are shrinking.

AN APPROACH TO DISTANCE EDUCATION BY USING NETWORK TECHNOLOGY

Dam Quang Hong Hai
University of Natural Sciences, HCMC, Vietnam

Abstract

Distance Education is now very popular in many countries, almost all foreign universities propose their distance courses. In Vietnam now, universities and colleges begin to set up the distance education. In this paper, we present one method for making distance course by using Network technology which can implemented in the condition of communication technique of Vietnam.

1. Introduction

With the economic development of the country, Vietnamese people pay more intention in education. A large number of people have demand to improve their knowledge by receiving the high education from the Universities or Colleges. Unfortunately, most of them live in far provinces, and the transportation between these provinces and big cities isn't convenient. To solve this difficulty, one solution for solving this difficulty is to make a distance learning courses, which can transmit from the center of education to the student's houses. The Universities of many countries provide their distance learning courses to public by using many media such as video tape, CD-ROM and computer networks like Internet or Intranet. But there is somehow different with Vietnamese environment.

2. Model of Distance Education system

To build a Vietnam Distance Education system, we must solve many difficulties of using Network technology. Some of them may be listed as follows:

- People just begin to use Internet .

- People can't use Internet facility to link with one server in site of Vietnam, they must use the long distance telephone line. If they use the long telephone facility, the price is very high compared with their incomes.

We provide the students a classroom atmosphere on which every student has the communication with the teacher and his friends. Some functions of our system are as follows :

- At first we built the teaching tools for the teachers. By using these tools, the teachers can make up the lesson with their familiar skill. In their lesson the teachers can include their writing speech from many word processing software, such as Word for Windows, Ventura, WordPerfect... and they can make the sample programs in Excel, C++

and include them in the lesson. These tools give the teacher more easy way to bring the same material for the in site students to the students of distance learning. When students receive the lesson they can read the text, test the sample in their computer.

- We build up a database of lecture knowledge, which help the teachers to control their work and use the difference sources of material to their lesson.

- After making the lesson, the teacher can transmit it to the students through the system of servers, which must be established in the University and some other provinces.

- All students who are enrolled in the course must have the student tools in their house and they can access one nearest server of the system by the telephone line. Through this link they can down load the lesson from the server to their computer , put in into the database and use it.

In this approach, we imitate the teaching process in which the teachers transmit their lecture with a lot of visual examples to their students. All this lecture note and those examples can be written in text and some another tools. All of those must be compressed into one file and passed from teacher's computer to student's. In addition, students can give the question to the teacher and receive the answers by the same way.

3. Implementation

When designing the system of distance education using the Computer Network we divide it into three parts , this means that different people need different part of the system.

- Teaching part is for the teacher , which consists of making lesson tool , database system of lessons , tool for reading the questions and answering the questions, communication part. In the teaching part, there are many tools that help the teachers to make their lesson. Lesson can contain the main text part , reference part, quiz part.

- Learning part is for the student , which consists of reading lesson tools , database system of lessons , tool for making the questions and reading the answers, communication part.

- Managing part is for the manager of the distance learning system, which is in the servers and consists of the lesson delivery system, question delivery system, answer delivery system, distance learning management system.

The relation between those parts are as follows.

- When the teachers want to make their lesson they use the teaching part in their computers to put all the information into the lesson then they use the communication part to send this lesson in compressed form to the University server.

- In the University server, the management part would pass this lesson to another servers, in which there are some students of this course and the management part on those severs would deliver this lesson to every student's box.

- When the students want to receive the lesson, they can connect to their server and receive all lessons , which are in their box. After receiving the lesson, the students can read every lesson by their computers as long as they want.

An important issue is how to divide the time for teacher and student. We divide their time into two parts.

- First part of time is the time for reading or writing the lesson .

- Second part of time is time for sending or receiving information .

We recognized that the time for teaching or learning is longer than the time for sending or receiving information . Based on this conclusion, the time for teaching or learning would be performed on computer of user and the communication link is established only on the time for sending or receiving information . Therefore, we can save time for communication

4.Conclusion

This approach satisfies the technical condition of Vietnam, We plan to move our system to the Internet environment. We hope, in the following years, when the price of using Internet services is reduced, we can use the Internet for our distance education system.

ABOUT THE WAYS TO SOLVE THE SHORTAGE OF IP ADDRESSES

Phan Cong Vinh
Vietnam Post and Telecoms Institute of Technology

Abstract

IP addresses are in high demand and short supply, but the shortage of IP addresses is no cause for warning. Net managers have several options for dealing with the queeze, such as ***First****, start with subnetting, in which a block of IP addresses assigned to a network is divided and spread out among separate, smaller networks.****Second****, consider private addressing--building a private network entirely out of unregistered IP addresses.* ***Third****, another possible fix (but one that's geared more toward ISPs and carriers) is classless inter-domain routing (CIDR), an address consolidation scheme that reduces the pressure on the Internet's core routers.* ***Finally****, there's IP version 6, a protocol upgrade that tackles the shortage head-on by expanding the address space from 32 to 128 bits--thereby vastly increasing the number of available addresses, and delivering enough IP addresses to last through the next millenium.*

1.Introduction

While Internet connectivity and intranets have become corporate networking must-haves, they're also providing businesses with a high-tech lesson in the laws of supply and demand because the IP addresses are in short supply . But the shortage in available IP addresses isn't about to bring the 'Net crashing down. The ISPs (Internet service providers), the InterNIC (the body that assigns IP addresses worldwide), and the IETF (Internet Engineering Task Force) are all addressing the address shortage, and they've come up with some solutions that net managers can put to work today.

- **First**, start with subnetting, in which a block of IP addresses assigned to a network is divided and spread out among separate, smaller networks.
- **Second**, consider private addressing--building a private network entirely out of unregistered IP addresses.
- **Third**, another possible fix (but one that's geared more toward ISPs and carriers) is classless inter-domain routing (CIDR), an address consolidation scheme that reduces the pressure on the Internet's core routers.
- **Finally**, there's IP version 6, a protocol upgrade that tackles the shortage head-on by expanding the address space from 32 to 128 bits--thereby vastly increasing the number of available addresses.

2.The first way : subnetting

The crisis wasn't always so acute. There once was a time when net managers seeking IP addresses could pretty much get what they asked for--typically a Class B network address supporting up to 65,534 nodes. With more than 16,000 Class Bs available, addresses were in plentiful supply.

The explosive growth of the Internet, along with the rise of intranets, has changed all that. Class B addresses are now harder to come by than ever before. Choosing a Class C address--which can handle 254 nodes--is an option, but most networks are larger than that, and cobbling together multiple Class C addresses isn't really the most elegant solution. In other words, there just aren't enough Class Bs left, while Class Cs just aren't enough.

Fortunately, there are ways to deal with the shortage in addressing . The simplest of these is subnetting--subdividing an IP network address to use it in several smaller networks.

Subnetting helps deal with one of the most glaring flaws in the present IP addressing system, which is that once a block of addresses has been assigned, all the host addresses in that block are forever consigned to that block. If some or all of them are never used (which is frequently the case), they're unavailable to anyone else.

Here's how subnetting works. Say an organization receives a Class B address of 172.16.0.0 (this is actually a reserved address, used here for illustration). The organization could split this address into up to 254 subnets by using addresses like 172.16.1.0, 172.16.2.0, and so forth. (In this example, the 0 is used for numbering hosts on that subnet.)

3..The second way : private network addressing

There's another way for net managers to get around the IP address crisis: use the special addresses that are reserved for private networks. IETF RFC (Request for Comment) 1918 sets aside three address blocks for use solely in private networks: Class A network 10.0.0.0, Class B networks 172.16.0.0 through 172.31.0.0, and Class C networks 192.168.0.0 through 192.168.255.0 (the RFC is available at http://www.ds.internic.net/rfc/rfc1918.txt). Originally, these reserved address blocks were intended for use in networks not connected to the Internet, or for isolated test or experimental networks. But the shortage has prompted networkers to use these blocks, hiding the private addresses behind firewalls or packet-filtering routers.

The obvious advantage of this scheme is that it makes the shortage of IP addresses a nonissue. But how these private addresses are translated into public addresses when they're sent to the Internet.

The key is NAT (network address translation), which is defined in RFC 1631 (see Figure 1). A firewall or router using NAT essentially takes all private addresses of outbound traffic (traffic from the internal network to the Internet) and converts the source address to that of the router or firewall's external interface (or to a series of addresses if there are multiple external interfaces). For inbound traffic, the process works in reverse: The NAT box converts destination addresses to those used by the private network. But address conversion is just one of the advantages of NAT. Security is another: Attackers can't go after machines they can't see--and private addresses aren't visible on the public Internet. Still, coming to bat with NAT means making some trade-offs. Using a firewall or router as a NAT box is a hard-and-fast requirement, and that means added cost, extra administration, and--perhaps--a performance penalty.

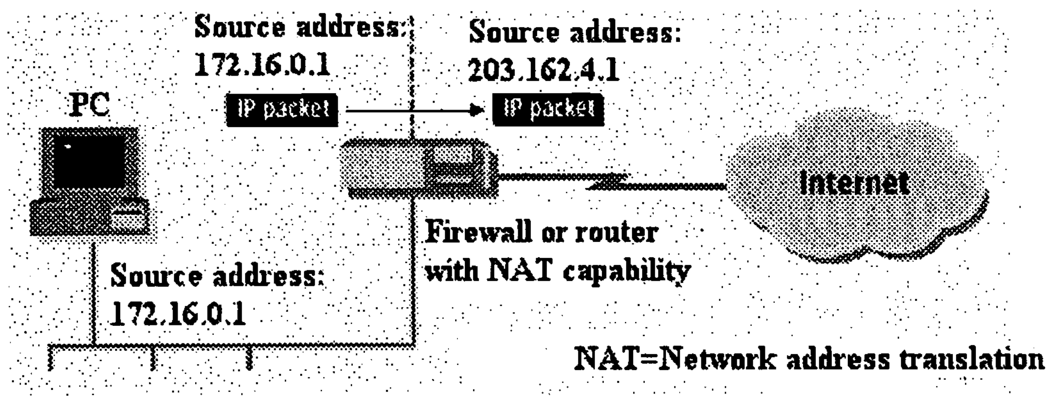

Figure 1

4.The third way: classless inter-domain routing

There's no question that private addressing is a good fix for the address shortage. The problem is that it's an option only for managers of private networks. ISPs face the problem of the addressing issue: keeping track of the huge amount of addresses being snapped up and put into use. National and international ISPs hook up with one another at network access points (NAPs). The routers at these Internet hubs have to know about every network on the Internet--unlike their counterparts lower down in the routing hierarchy, which have to know about just a few networks and can point to default gateways for the thousands they know nothing of. The NAP routers have no such luxury: They are the default gateways. Further, each new IP network added to the Internet requires a new NAP routing table entry. As more and more entries are made, the routing tables may become too large; parts of the 'Net then begin to fall off, rendering those networks unreachable.

That's prompted ISPs to turn to CIDR. Although it's not really a solution to the IP address shortage itself, CIDR reduces the number of routing table entries by consolidating addresses into contiguous blocks. If a range of addresses belongs to one ISP, the routers have to know only the range of addresses served by that ISP, not the individual network addresses. And when NAP routers have fewer table entries, they're likely to perform better and be able to see all the networks attached to the Internet. CIDR, described in

RFCs 1517 through 1520, is "classless" addressing. It replaces Class A, B, and C addresses with a network number "prefix" and a "mask" .Together, the prefix and mask identify a block of IP network numbers. All the addresses within a CIDR block are served by a specific ISP as part of a so-called autonomous system (AS), which usually means the group of routers belonging to that ISP. The ASs use the border gateway protocol (BGP) to exchange routing information with one another (see Figure 2). Within each AS, routers update one another using the same routing protocols they've always used, whether RIP (routing information protocol), IGRP (interior gateway routing protocol), or OSPF (open shortest path first).

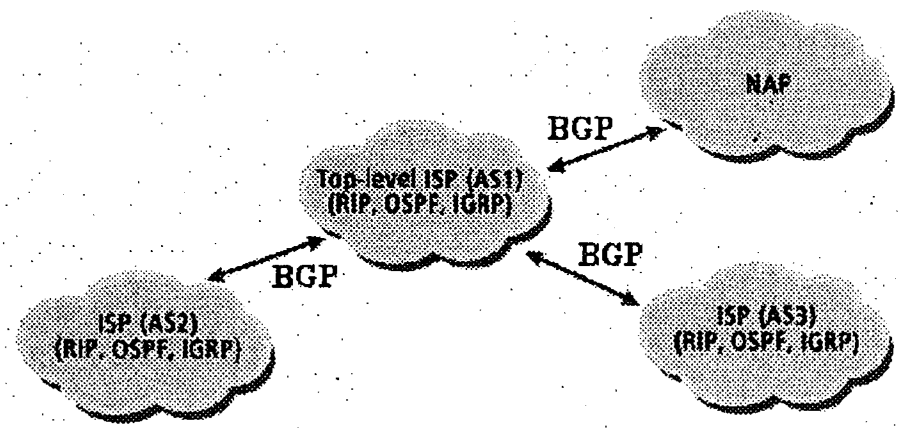

BGP=Border gateway protocol **NAP=Network access point**
CIDR=Classless inter-domain routing **OSPF=Open shortest path first**
IGRP=Interior gateway routing protocol **RIP=Routing information protocol**
ISP=Internet service provider

Figure 2

So, how does that all result in the reduction of routing table size? Consider an ISP that services 254 Class C network addresses, starting with 204.36.0.0. The addresses start with network 204.36.1.0 and run to 204.36.255.0. The CIDR notation for all of the networks in this block is 204.36.0.0 /16, where "/16" is the CIDR mask. The first 16 bits of the 32-bit IP address--204.36--identify the starting network number of the CIDR block. The remaining bits identify what were formerly considered separate Class C networks.

Eliminating class distinctions gives ISPs more flexibility in handing out addresses. For example, an ISP could elect to subdivide the /16 CIDR block into two /17 CIDR blocks, each with 128 contiguous networks, or into four /18 CIDR blocks, each with 64 contiguous networks. Note that adding a bit to the CIDR mask reduces by a power of two the number of contiguous networks in the block--254 networks with a /16 mask, 128 networks with a /17 mask, and so on.

Regardless, the ISP has just one routing table entry as far as the top-level NAP routers are concerned. There's no longer any need to know exactly where all of the networks in the 204.36.0.0 address block are located. So when a NAP router sees an IP datagram bound for any address that starts with 204.36.x.x, it locates the single routing table entry for that CIDR block. Then it tosses the packet on the doorstep of the ISP whose AS owns the addresses--leaving delivery up to the provider.Sprint Corp. (Kansas City, Mo.) and some other large ISPs have called for a minimum CIDR block size. Their goal is to restrict the number of routing table entries by forcing customers (usually other ISPs) to aggregate a minimum number of addresses into CIDR blocks. So far (to allow for an orderly transition), some top-level ISPs are doing this on newly assigned IP address space only, starting with /18 CIDR blocks in the 206.0.0.0 address block. Sprint is grandfathering in older addresses, but it and other members of the North American Network Operators Group (NANOG) are pressuring ISPs to use CIDR for previously assigned addresses, too.With the /18 prefix set out as a stipulation, an ISP can announce to Sprint a block of 64 network addresses, but not a smaller block of 32 (/19)--which would be filtered out by Sprint's routers. In other words, Sprint won't know where networks in those smaller CIDR blocks are (because its routers won't list them), which means systems behind Sprint's network won't be able to reach them.

CIDR may not solve IP address exhaustion, but when it comes to allocating the right number of addresses the scheme is a big help.Say a network manager needs network addresses for 10,000 hosts. That normally means applying for a Class B address--a request likely to be denied given how scarce Class Bs are. Even if a Class B were granted, more than 55,000 addresses would go unused (remember, Class Bs support more than 65,000 hosts). But with CIDR, a net manager can apply to an ISP for a block of 64 Class Cs. The CIDR scheme offers plenty of room for growth--a block of 64 Class C addresses supports more than 16,000 hosts--without unduly draining the pool of available addresses. That doesn't mean there are no CIDR downsides. ISPs, for instance, will look to serve only those addresses within their CIDR blocks. A network whose addresses are outside that block might be dropped from some routing tables, cutting it off from other parts of the Internet. Network managers could renumber networks with another address that is part of a CIDR block, but doing so tends to be costly and time-consuming. Still, as long as the public network address used by their routers or firewalls is part of a CIDR block, net managers are unlikely to feel the effects of the CIDR scheme. (There is a method for determining how a network address is announced by core routers. Instructions are available at http://www.ra.net/RADB.tools.docs/.query.html.)

5.The final way: IPv6

But even if CIDR addressing is fully implemented by every ISP in the world, the addresses will simply one day run out. It's inevitable. It's a problem that the developers of IP only dimly foresaw 20 years ago, but now the IETF is moving to counteract the shortage. It has sanctioned an upgrade known as IP version 6 (IPv6), which dramatically expands the number of available addresses by boosting them from 32 to 128 bits.What's more, IPv6 will, through the use of a hierarchical routing scheme, ease the workload of routers. CIDR blocks can be aggregated on the basis of geographical location or ISP assignment--which enables routers to determine where a network is located by its address. That's a big change from IPv4, under which ISPs and the InterNIC assign addresses at random. Behind this scheme lie new exterior routing protocols--such as the OSI interdomain routing protocol (IDRP)--that promise to improve router performance by carrying CIDR masks as well as IP addresses. The main idea with hierarchical routing is that a site's networks are part of a small CIDR block, which is part of a larger CIDR block from an ISP, which in turn is part of a regional or continental CIDR block. Routers in other regions carry the largest CIDR blocks in their routing tables and use them to forward traffic for any network in the block to NAPs in appropriate locations. However, even with bigger CIDR blocks, there will still be more and more networks--which means bigger, more powerful, and faster routers will be needed.

6.Conclusion

On paper, IPv6 is a great idea. It will relieve the IP address crunch, and it promises to streamline configuration and management of workstations and routers alike. Still, IPv6 poses some daunting questions for net managers. For instance, what's the best way to make the transition while maintaining backward compatibility with all those systems still running IPv4? What about renumbering networks--not to mention buying, installing, and configuring all of that new IP software?

IPv6 advocates say there's no reason to worry. The transition plan calls for 32-bit IPv4 addresses to be embedded in the least-significant-bit positions of the IPv6 address field. This would permit communication between IPv4 and IPv6 systems and would allow a system to run dual protocol stacks until the day IPv4 is officially replaced.

The downside is that this plays into the hands of networkers who don't want to convert to IPv6 at all; instead, they have a workable option for continuing to run IPv4. For network managers who do want to make the move, about the only thing they can do right now is make sure their networks are part of CIDR address blocks. Ideally, all addresses should be part of one contiguous block--but that may not be possible for enterprise nets.

In short, the switch to IPv6 won't come about until the value of a new technology becomes clear and system hardware and software can support it. Today, only a handful of vendors offer production-grade IPv6 products. Net managers will thus continue to wring as much out of IPv4 as they can--until their own systems or the sheer numbers of Internet users forces the conversion to IPv6.

Reference

[1] IETF RFC 1918 *http://www.ds.internic.net/rfc/rfc1918.txt*

[2] IETF RFC 1517 *http://www.ds.internic.net/rfc/rfc1517.txt*

[3] IETF RFC 1518 *http://www.ds.internic.net/rfc/rfc1518.txt*

[4] IETF RFC 1519 *http://www.ds.internic.net/rfc/rfc1519.txt*

[5] IETF RFC 1520 *http://www.ds.internic.net/rfc/rfc1520.txt*

[6] Gilbert Held, IP for the Next Generation, July-1, 1997.
http://www.networkmagazine.com/9707/9707ip.htm

[7] The information from *http://www.ra.net/RADB.tools.docs/.query.html*

INTRODUCTION TO VERY LARGE DATABASES

Do Hoang Cuong
University of Natural Sciences, HCMC, Vietnam

Abstract

When administering a data warehouse or Very Large Database (VLDB) environment, a number of items must be considered and re-examined in light of the special issues facing a VLDB. The main challenge of a VLDB is that everything is larger and maintenance task take considerably longer. VLDBs provide unique challenges that sometimes require unique solutions to administer them effectively.

Contents

1. What is a VLDB?

Some people, it depends on how you define "very large":

- A fixed size (for example 2G, 20G, 200G, 400G ...) ?
- When database-restore time exceeds a certain threshold.

In general, there's no easy way to quantify the point when a database becomes a VLDB. Here's the preferred definition :

A very large database is any database in which standard administrative procedures or design criteria fail to meet business needs, due to the scale of the data.

In other words, whenever the size of the database causes you to redesign your maintenance procedures or redesign your database to meet maintenance requirements, you're dealing with a VLDB.

2. VLDB Maintenance Issues

VLDBs including these important ones :

1. Time required to perform dumps and loads.
2. Time required to perform necessary database-consistency checks.
3. Time and effort required to maintain data.
4. Managing partitioned databases.

3. Explores These Issues And Provide Guidelines For Implementing Solutions

3.1. Managing Database Dumps and Loads

- Database dumps are necessary to provide recoverability in case of disaster. The main issue with database dumps and VLDBs is the duration of the database dumps; dumps time is proportional to the amount of data in the database.

- If the time to back up a database is excessive, the time to restore it is even more. The ratio between dump and load duration is approximately 1:3.

- VLDB backup/restore procedure :

1. Consider the amount of time you're willing to be "down" while performing a recovery. If you need a database to be recovered within 8 hours, for example, determine the size of a database that can be recovered in that amount of time.
2. Estimate table sizes for your database, to determine partitioning sizes and options:
   - If you have a 40G logical database, for example, you may need to implement ten 4G databases.
   - Are any single tables greater than 4G ?
3. How many tables can you fit in a database ?
   - Are any tables candidates for partitioning, based on this determination alone ?
4. Develop your administration schedule:
   - For every day during a month, determine what periods of time can be dedicated to administrative activities.
   - Are weekends available ?
   - If you determine that you have five hours per night to perform administrative activities, you then need to determine what activities need to be completed and how they can be distributed over your available administrative time.
5. Determine the rate of the dump process
6. Finalize your schedule and document accordingly.
7. Monitor and update the process as needed over time as thing change.

3.2. Checking Database Consistency

Almost Database Server (SQL, ORACLE, DB2,...) provide a systems administration tool that verifies pointers internal to a database and its structures. Remember, this tool should be run prior to any database dump to avoid dumping a corrupt database. The worst time to realize you have a bad page pointer is during recovery of critical database when the load process fails due to inconsistency.

This tool typically lock user tables, indexes, system tables and databases when running. This tool is very I/O intensive. The more I/O required, the longer it takes. These are the main issues with running this tool in VLDBs.

To develop a plan for effectively checking your database consistency in a VLDB, you first need to analyze your tables and rank them in order of importance. For example, where would a corruption have the most serious effects on your system ?

Next, analyze each non-clustered index on your table and rank it in order of importance to determine which indexes are most important to check.

Plan to check your high-activity tables as close to your dump as possible, because thes tables are more likely to encounter allocation problems. Verify the consistency of these tables as close to the dump as possible. Static tables should need to be checked only after data loads.

3.3. Data Maintenance

To performing database dumps and checking database consistency on a VLDB, Usually there are other data-maintenance tools that need to be performed on a database. These include purging and archiving data.

The data in VLDB may grow to a size approaching or exceeding the maximum available database space. At this point, the decision needs to be made either to expand the database or to purge or archive data to free space. Purging of data is often necessary. When purging or archiving data, a number of issues need to be addressed:

1. Locking

One way to avoid locking problems is to perform archival activities when users aren't using the system. Here are alternatives to avoid table-level locks:

- Use cursors to restrict rows being modified to one at a time.
- Use set row-count to affect a limited number of rows at a time.
- Use some other table specific value to limit the number of pages that are being modified to less than the lock-escalation threshold.

2. Logging

Rows being deleted from a table are logged in the transaction log. You determine whether your transaction log is large enough to handle the deletion of a large number of rows in a single transaction. To minimize I/O contention, your log should be placed on a separate disk to distribute the I/O.

Although your transaction log may be large enough to handle a single deletion of 500000 rows from a table, those records remain in the log until they' re dumped. Your purge/archive process should be designed to dump the transaction log after the completion of the purge process, to clear the log for the next purge or normal system activity.

If your log isn't large enough to handle the deletion as a single transaction, break the deletion into multiple transaction, dumping the logs between each transaction.

3. Referential Integrity (RI)

The following are items to consider when dealing with RI:

- If referential integrity is maintained via triggers, cascading delete triggers may exist on a table where data is to be removed. The deletion of rows from one table may cause the deletion of even more rows from a related table. Even though your log may be able to handle the delete from the first table, the cascading defects could fill your log, causing the purge to fail. It could also lead to exclusive table locks throughout you system.

- If the purge/archive process is a batch job run during administrative periods, you can drop the triggers or RI constraints before running the purge/archive. If the purge/archive is to be conducted during business hours, you probably should leave the triggers on the table.
- If there are cascading delete triggers, you may need to perform deletes in small quantities with regard to the number of rows deleted in each table in the relationship, to avoid table-level locks.

4. Transactional integrity

A transaction is a logical unit of work. All activities in the transaction must complete successfully or they all should fail. During an archive process, you'll likely insert data into a table in different database and remove the data from its current location. If an error occurs that prevents the process from inserting the data into the new database, your archive process mustn't delete the data from its current database. Therefore, the insert/delete activity should be conducted in a transaction.

When archiving data with relationships across multiple tables, the design may require you to write transactions that transfer the data from all the related tables before deleting any of the rows. Make sure that your design considers multi-table relationships.

3.4. Data-Partitioning

- When dealing with VLDBs, it may become necessary to partition the database, due to Database Server size limitation or in order to meet your backup and recovery or data-maintenance requirements. There are two primary ways of partitioning databases: vertical partitioning and horizontal partitioning.

- Vertical partitioning of data is the process of drawing imaginary lines through a database schema and placing individual tables in different databases. It may also involve partitioning individual tables by columns, to separate the frequently accessed columns from infrequently accessed columns, and then placing the resulting tables in the same or different databases.

- Horizontal partitioning of data involves breaking up tables into logical subsets and placing them into the same or different databases.

References

[1] Microsoft SQL Server 6.5 DBA Survival Guide- SAM Publishing - Orryn Sledge & Mark Spenik.

[2] DB2 - Administration Getting Started -IBM Universal DB2 For Windows NT Book On-Line.

[3] ORACLE Unleashed - SAM Publishing - Rachel Becker, Mathew Bennett, etc...

U.S. DEPARTMENT OF EDUCATION

OFFICE OF EDUCATIONAL RESEARCH AND IMPROVEMENT (OERI)

EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

REPRODUCTION RELEASE (Specific Document)

I. DOCUMENT IDENTIFICATION

Title: IT@EDU98: INFORMATION TECHNOLOGY IN EDUCATION TRAINING. PROCEEDINGS. JANUARY 15-16, 1998, HO CHI MINH CITY, VIETNAM

Author(s): Dr. Thomas Owens and others

Corporate Source (if appropriate): Northwest Regional Educational Laboratory

Publication Date: 1998

II. REPRODUCTION RELEASE

In order to disseminate as widely as possible timely and significant materials of interest to the educational community, documents announced in the monthly abstract journal of the ERIC system, Resources in Education (RIE), are usually made available to users in microfiche and paper copy (or microfiche only) and sold through the ERIC Document Reproduction Service (EDRS). Credit is given to the source of each document, and, if reproduction release is granted, one of the following notices is affixed to the document.

If permission is granted to reproduce the identified document, please CHECK ONE of the options and sign the release below.

CHECK HERE ▶ [XX] Microfiche (4" x 6" film) and paper copy (8½" x 11") reproduction

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY Northwest Regional Educational Laboratory TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

OR [ ] Microfiche (4" x 6" film) reproduction only

"PERMISSION TO REPRODUCE THIS MATERIAL IN MICROFICHE ONLY HAS BEEN GRANTED BY ____ TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

Documents will be processed as indicated provided reproduction quality permits. If permission to reproduce is granted, but neither box is checked, documents will be processed in both microfiche and paper copy.

SIGN HERE ▶

"I hereby grant to the Educational Resources Information Center (ERIC) nonexclusive permission to reproduce this document as indicated above. Reproduction from the ERIC microfiche by persons other than ERIC employees and its system contractors requires permission from the copyright holder. Exception is made for non-profit reproduction of microfiche by libraries and other service agencies to satisfy information needs of educators in response to discrete inquiries."

Signature: Jerry D. Kirkpatrick

Printed Name: Jerry D. Kirkpatrick

Organization: Director, Institutional Development and Communications, Northwest Regional Educational Laboratory

Position:

Address: 101 S.W. Main St., Suite 500, Portland, OR

Zip Code: 97204

Tel. No.: (503) 275-9517

Date: 3/13/98

III. DOCUMENT AVAILABILITY INFORMATION (Non-ERIC Source)

If permission to reproduce is not granted to ERIC, or, if you wish ERIC to cite the availability of the document from another source, please provide the following information regarding the availability of the document. (ERIC will not announce a document unless it is publicly available, and a dependable source can be specified. Contributors should also be aware that ERIC selection criteria are significantly more stringent for documents which cannot be made available through EDRS.)

Publisher/Distributor:

Address:

Price Per Copy:

Quantity Price:

IV. REFERRAL TO COPYRIGHT/REPRODUCTION RIGHTS HOLDER

If the right to grant reproduction release is held by someone other than the addressee, please provide the appropriate name and address: