



UNIVERSITY OF MINNESOTA



UNIVERSITY OF MINNESOTA

A Place for your Stuff

Rolling Out Large-Scale iSCSI at the UofM

Christopher R. Hertel

University Data Management Services

August, 2005

We Are Here



U of M Overview

- Big Ten University
- About 61,000 students
- Over 48,000 in the Twin Cities alone
- “600 Organizations with One Name”

The Mission

- ➔ Improve Productivity
- ➔ Cut Costs

The Method

- ➔ Provide Consolidated Storage for:
 - ✓ Efficiency
 - ✓ Flexibility
 - ✓ Supportability

The Madness:

A Campus-wide SAN

Sanity Check:

- ◆ Roll out services in phases
- ◆ Data-center services first

UDMS started by building out a campus FibreChannel network connecting major datacenters.

- St. Paul
- East Bank
- West Bank

This provided an opportunity to consider next steps before taking them.

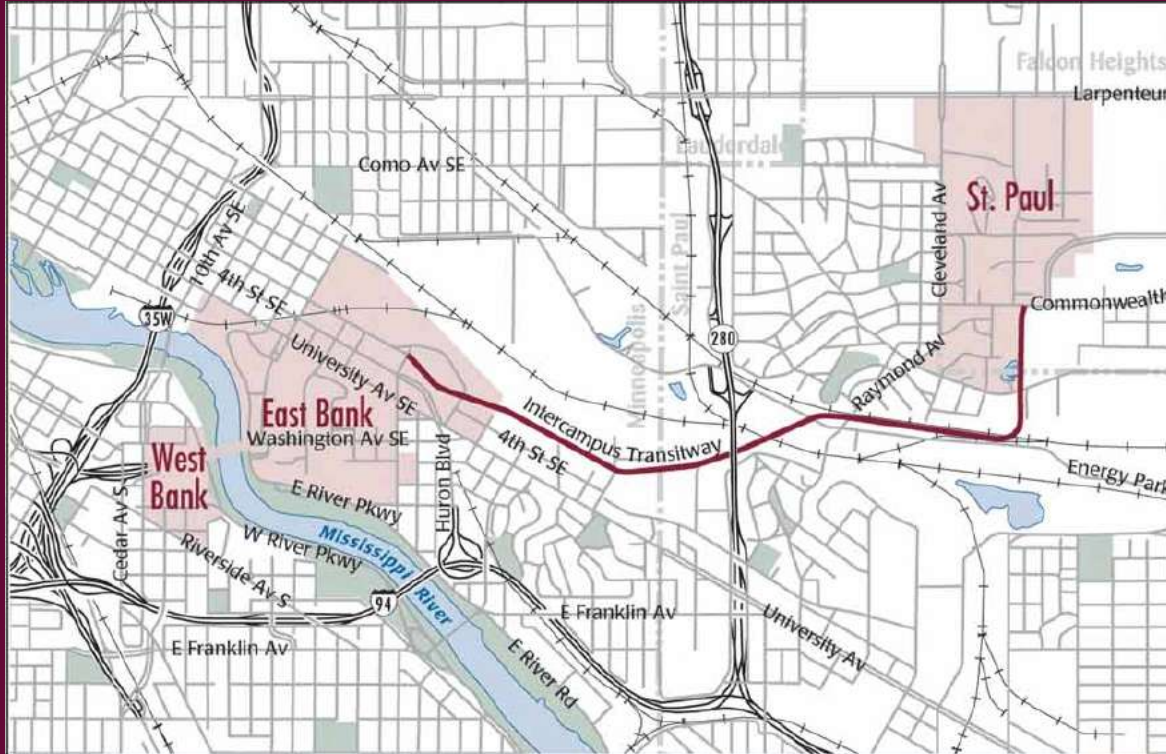
Trouble Ahead

 FC Network Limitations

 Space Constraints

 Costs

Campus Networking Challenges



- Mississippi River
- Distance from Minneapolis to St. Paul

FibreChannel Deployment Limits

- Physical Space
- Available Fiber
- Cost to overcome these problems

Space and Fiber are carefully guarded resources.

As a result:

- FibreChannel is the high-end solution for enterprise-class systems.
- The datacenters have become SAN Prisons.

iSCSI is the Key to Unlock the SAN Prison

...or so we've been told,

but we had concerns:

- ✗ Few comparable installations
- ✗ Routed networks were an unknown
- ✗ Impact on Campus Network was unknown
- ✗ The Hardware/Software Initiator debate

iSCSI Testing: Goals

Ensure that iSCSI would work in our diverse academic environment.

④ Reduce the number of products (variables) by elimination through testing

④ Support the maximum variety of clients

④ Provide “enterprise-level” service

Hardware HBA's seemed the logical choice.

iSCSI Testing:

The “Get Comfy” Phase

Kick the tires, and take iSCSI for a test drive.

- 🐱 Install hardware HBAs
- 🐱 Present LUNs via iSCSI
- 🐱 See how it goes...

We also tried software Initiators, just because we were curious about them.

Discoveries

- ❁ iSCSI is resilient
- ❁ Applications are less resilient
- ❁ iSCSI \Leftrightarrow FC bridges don't hide FC complexity
- ❁ iSCSI “servers” do hide FC complexity
- ❁ Not enough drivers for hardware HBAs

iSCSI Testing: Performance

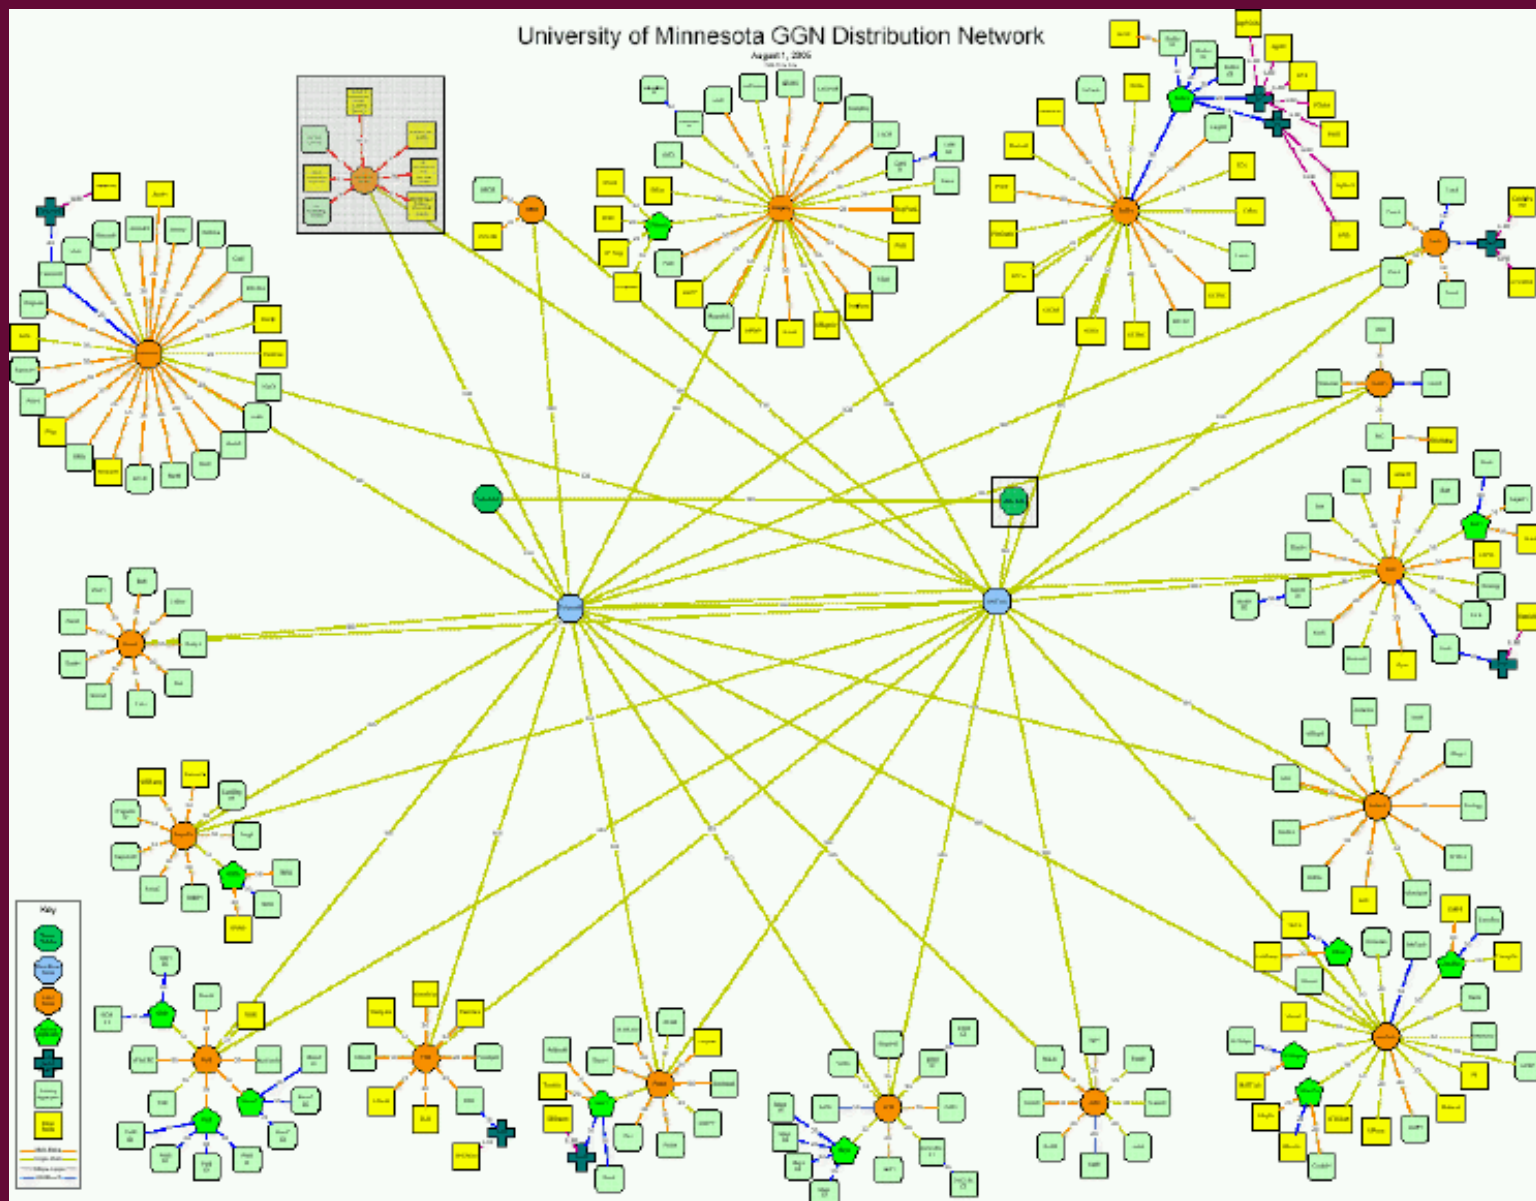
- ❖ At low speeds, throughput is a reasonable fraction of available bandwidth
- ❖ Faster than tested NAS protocols (in simple file copy tests)
- ❖ Software initiators are as fast as (faster than?) hardware initiators

Looks as though we'll be using software initiators.

Next Step:

Dealing with the Campus Net

- ▶ Avoiding the Firehose phenomenon
(being good net citizens)
- ▶ Security Concerns
- ▶ “Rightsizing” the service



Firehose Phenomenon

The Gopher Giganet is designed to support many servers, distributed across campus.

- Concentrating “central” services to one location is probably a bad design.
- iSCSI traffic might impact other datacenter services
- Other services might impact iSCSI

Limited testing means that we do not know if these are realistic concerns.

Firehose Phenomenon

Solution: Avoid the problem by serving iSCSI from multiple datacenters.

- ⌚ Spreads the load
- ⌚ Reduces impact of localized network outages
- ⌚ Aids in problem isolation
- ⌚ Might be used to provide redundancy

Security Concerns

Some customers want to store sensitive data:

- ☺ FIRPA
- ☺ HIPAA
- ☺ Juicy E'mail

Not a simple problem...

- Virtual Routing & Forwarding to logically isolate iSCSI traffic
- IPSec for desk-to-datacenter encryption
- Encrypted filesystems for end-to-end coverage

Rightsizing

Initial “enterprise service” assumptions have been proven wrong. Do we need to rethink the assumed customer profile?

“A Place for you Stuff”

Is the campus SAN a “Data Warehouse”, or a “Digital MiniStorage”?

Rightsizing

Potential uses include:

- ✿ Enterprise application servers w/o FC access
- ✿ File servers that just need more disk
- ✿ Desktops doing research data collection
- ✿ My boss, who just wants to back up his Mac

Other Stuff to Think About

iSCSI is one-to-one

- Can provide disk to NAS
- Some have inquired about clustering

Local storage is cheap

- We can offer better service & support
- Departments will set up their own anyway

Beyond iSCSI

Someone noticed that our iSCSI server can also serve NFS and CIFS

- ⊗ Except for the NDS tree, U NAS services are typically departmental
- ⊗ Deploying centralized NFS and CIFS amplifies the Firehose
- ⊗ CIFS requires domain administration
- ⊗ Needs more “Get Comfy” testing

Beyond iSCSI

The “One Gig per User” Initiative

- ❖ Delivering a “virtual minidisk”
- ❖ Accessible on campus and off
- ❖ Cross platform
- ❖ Collaboration Zones
- ❖ Expandable Feature set

Looks like a job for WebDAV...
...maybe.

The End