

IBM Lotus® Sametime® Large-Scale Architecture Design

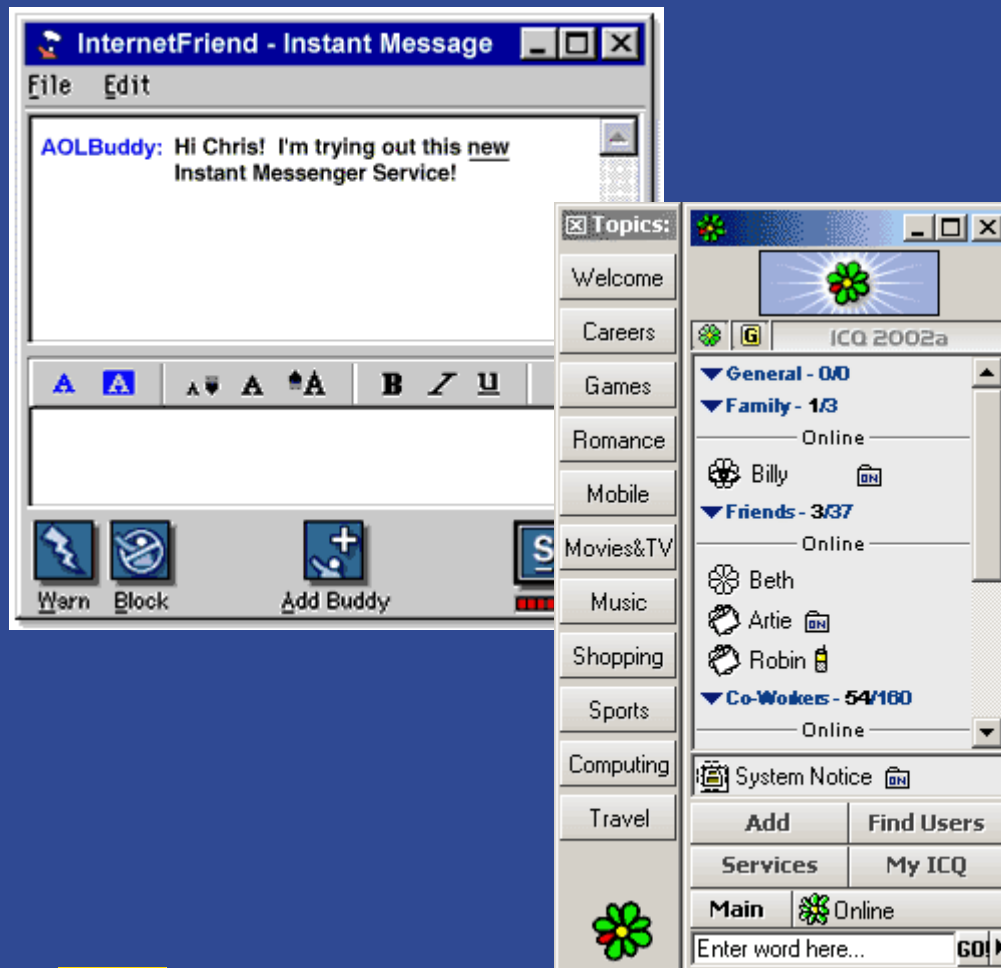
Jason Roy Gary, Executive Architect, IBM Greater China Group



What is in this Session?

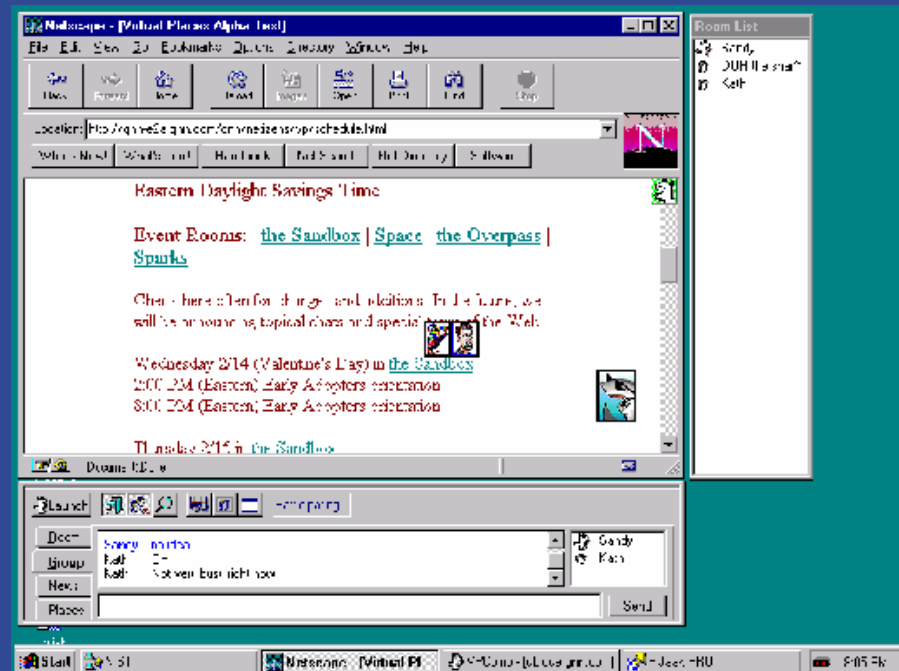
- A Historical Perspective
 - ▶ Large Scale Instant Messaging
- A Sametime Architecture Review
- Lotus Sametime Customer Based Testing Results
- Basic Architectural Principals
- Real World Architectural Examples

A Historical Perspective



Instant Messaging – A Quick History

- Before Instant Messaging
 - ▶ CompuServe® – Private Rooms
 - ▶ BBS – One Modem at a Time
 - ▶ IBM AIX® (Hey, etc.)
 - ▶ Ntalk / YTalk
- Q-Link® (1985)
- Ubique acquired by AOL 1995
- Buddy List, AIM® – 1996
- ICQ® (I Seek You) – 1996 (purchased by AOL, 1998)
- Online Group Chat (ICQ, 1998)
- Ubique acquired by IBM 1998
- Lotus Sametime 1.0 December, 1998
- Instant Meetings – Lotus Sametime, 1999
- Federated Buddy Lists (Trillian®, 2002)
- Jabber (2002)
- LCS 2003 (2003)
- Yahoo, AOL Leave Enterprise Messaging Space (July, 2004)
- Lotus Sametime 7.0 / LCS 2005 (2005)
- Lotus Sametime 7.5 (2006)



Virtual Places – A Quick History

- The protocol that enables Lotus Sametime presence is called “Virtual Places”
- It was first launched in 1994 and even included basic voice capabilities
- It was originally intended to allow users to chat in real time about games, or web pages
- At one time it was the IM system for Excite and Netscape
- VP was used by AOL to hold what was then the largest web collaboration meeting, in 1995
- It still lives on, reverse engineered, on many sites, especially gaming sites
- Today Lotus Sametime uses an evolution of VP

Maturity vs. Immaturity

- Instant Messaging and Presence as we know them have only existed since 1996 – 11 years
- By contrast the first “internet” @ email system was created in 1971
- T120 was not ratified until 1993 – 1995 for web conferencing
- Instant Messaging and Web Conferencing are new and immature technologies that are still growing
- For example originally VP was not designed to support more than one presence server in a community and could not do so until 1998 – for that matter a multi-server IM system did not exist in production until 1999!!
- Most of the architectural challenges in scaling Lotus Sametime are a function of the historical design of the product AND the immaturity of the services of IM and Web Conferencing

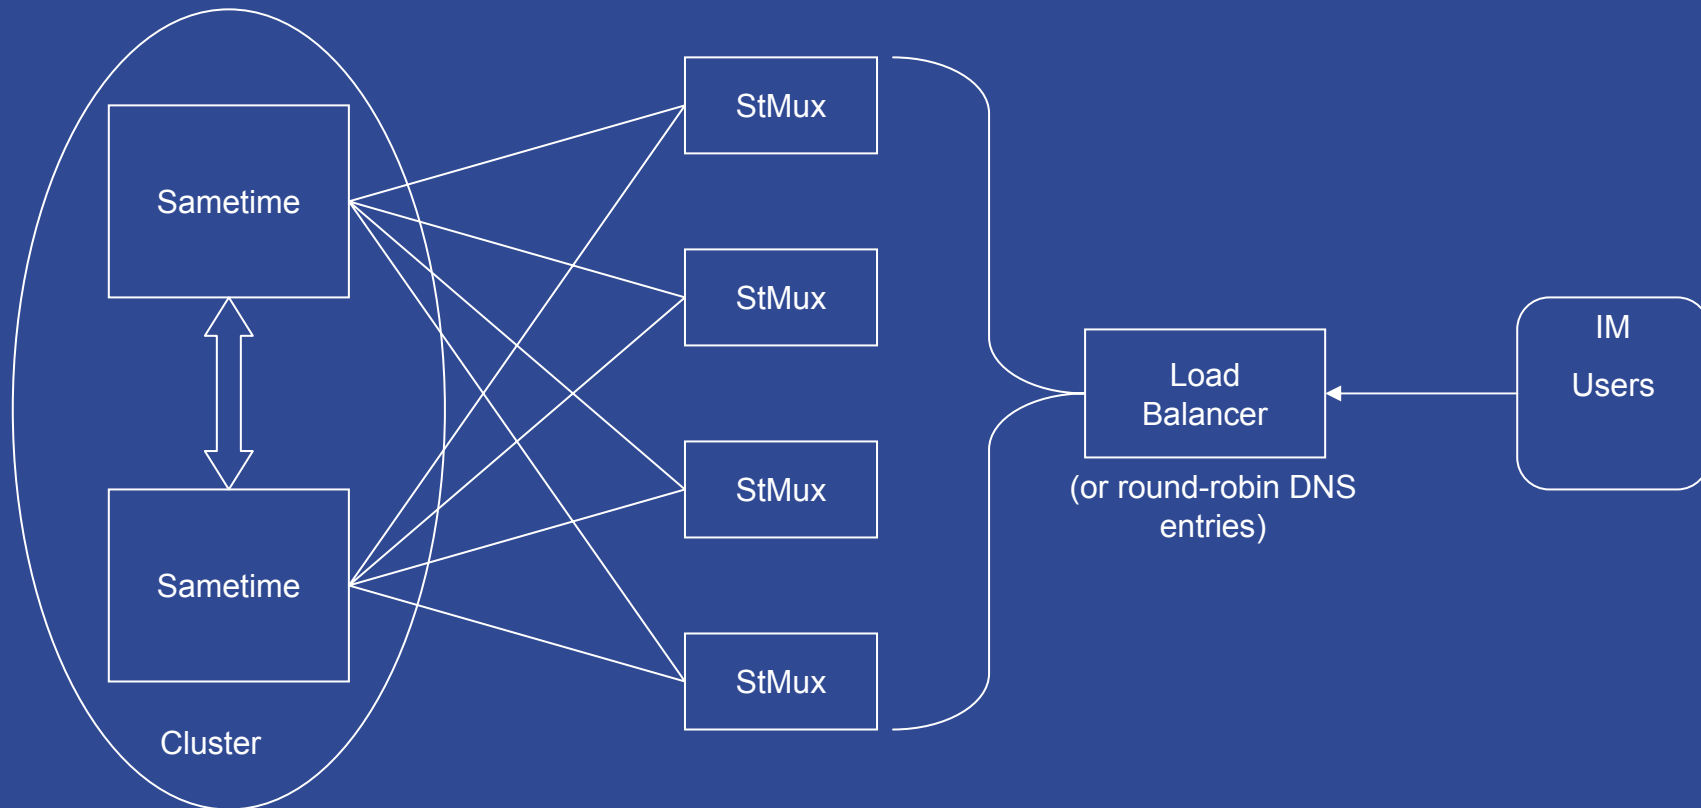
A Basic Review



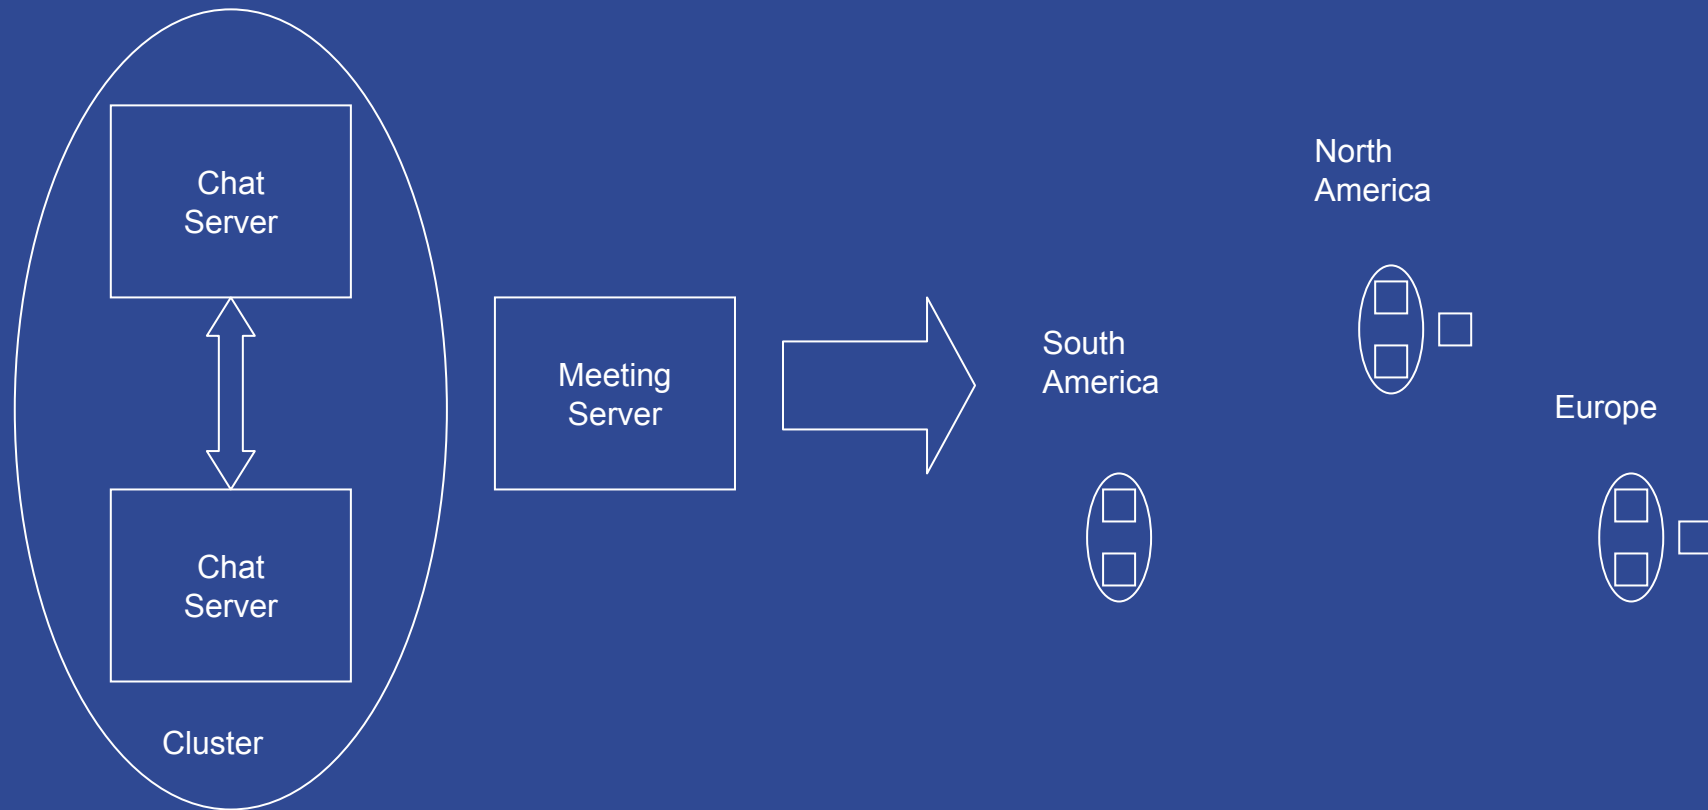
Basic Review

- The Lotus Sametime Community is defined by its directory
- Directory choice – IBM Lotus Domino[®] Address Book or LDAP
 - ▶ One or the other
- The home server is the base of one's awareness in the community
 - ▶ Buddy list storage
 - ▶ Presence logical "home"
- Users can log into any server, but:
 - ▶ Users are "proxy" connected to their home server
 - ▶ Awareness is maintained on home server
 - ▶ Buddy List retrieved/managed on home server
 - ▶ Instant Meetings launch on the home server of initiator - ALWAYS
- Map users to home servers by geography (or network), NOT job function!
- Buddy lists are stored in a Notes database, vuserinfo.nsf (with one exception)

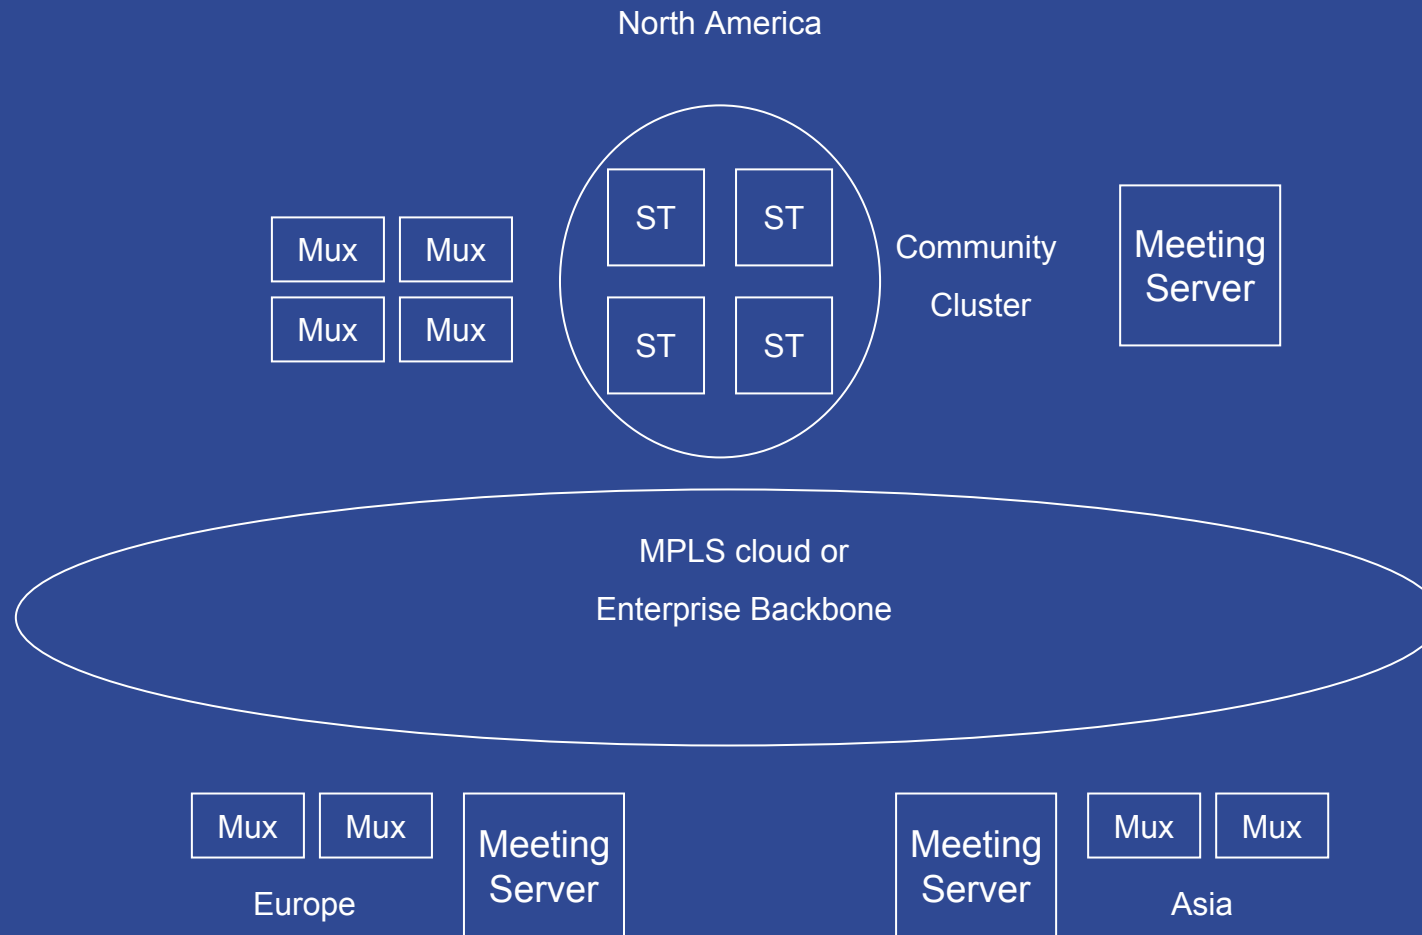
Deployment Architecture: Community Clusters



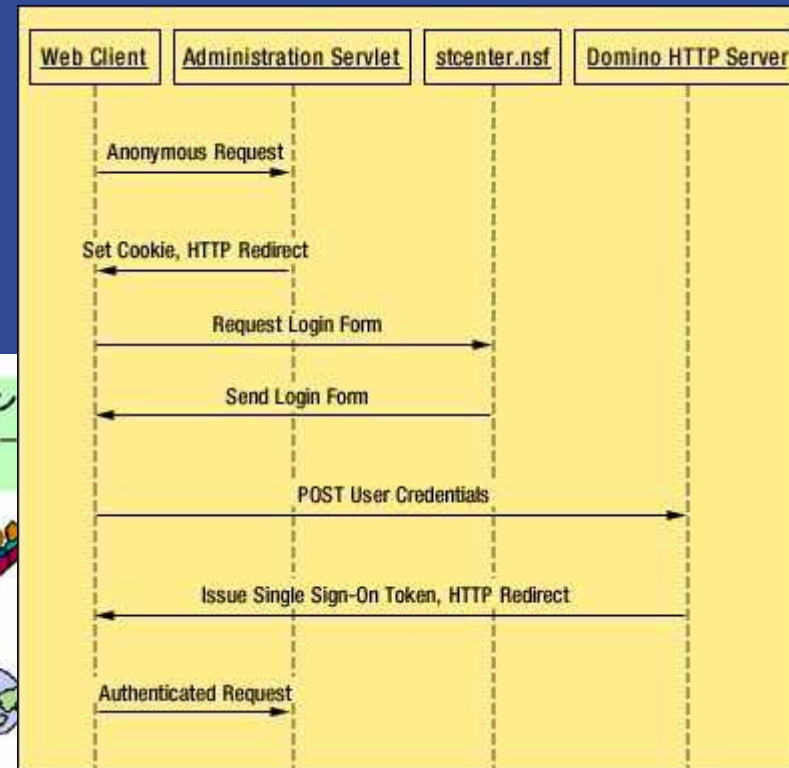
Deployment Architecture: Splitting Services



Deployment Architecture: Combining Techniques



Lotus Sametime Real World Testing Results




Lotus Sametime Real World Testing Parameters

- Performed both in IBM's Lab with Customer Data and On-Site at Customer on the Customer's hardware and network
- ST1 / ST2 denotes a cluster of two community servers, two MUX servers and, two LDAP servers (see figure to right)
- Servers were running Windows 2003
- All servers were quad processor servers with 4GB of RAM
- These tests were with Lotus Sametime 6.5.1 FP1 on Lotus Domino 6.5.4 FP1
- Directory size was > 400,000 users and > 90,000 groups
- LDAP Servers were running Lotus Domino 6.5.4
- Logins were only limited by the number of testing "driver" machines

Lotus software

ST1 / ST2
Testing Configuration



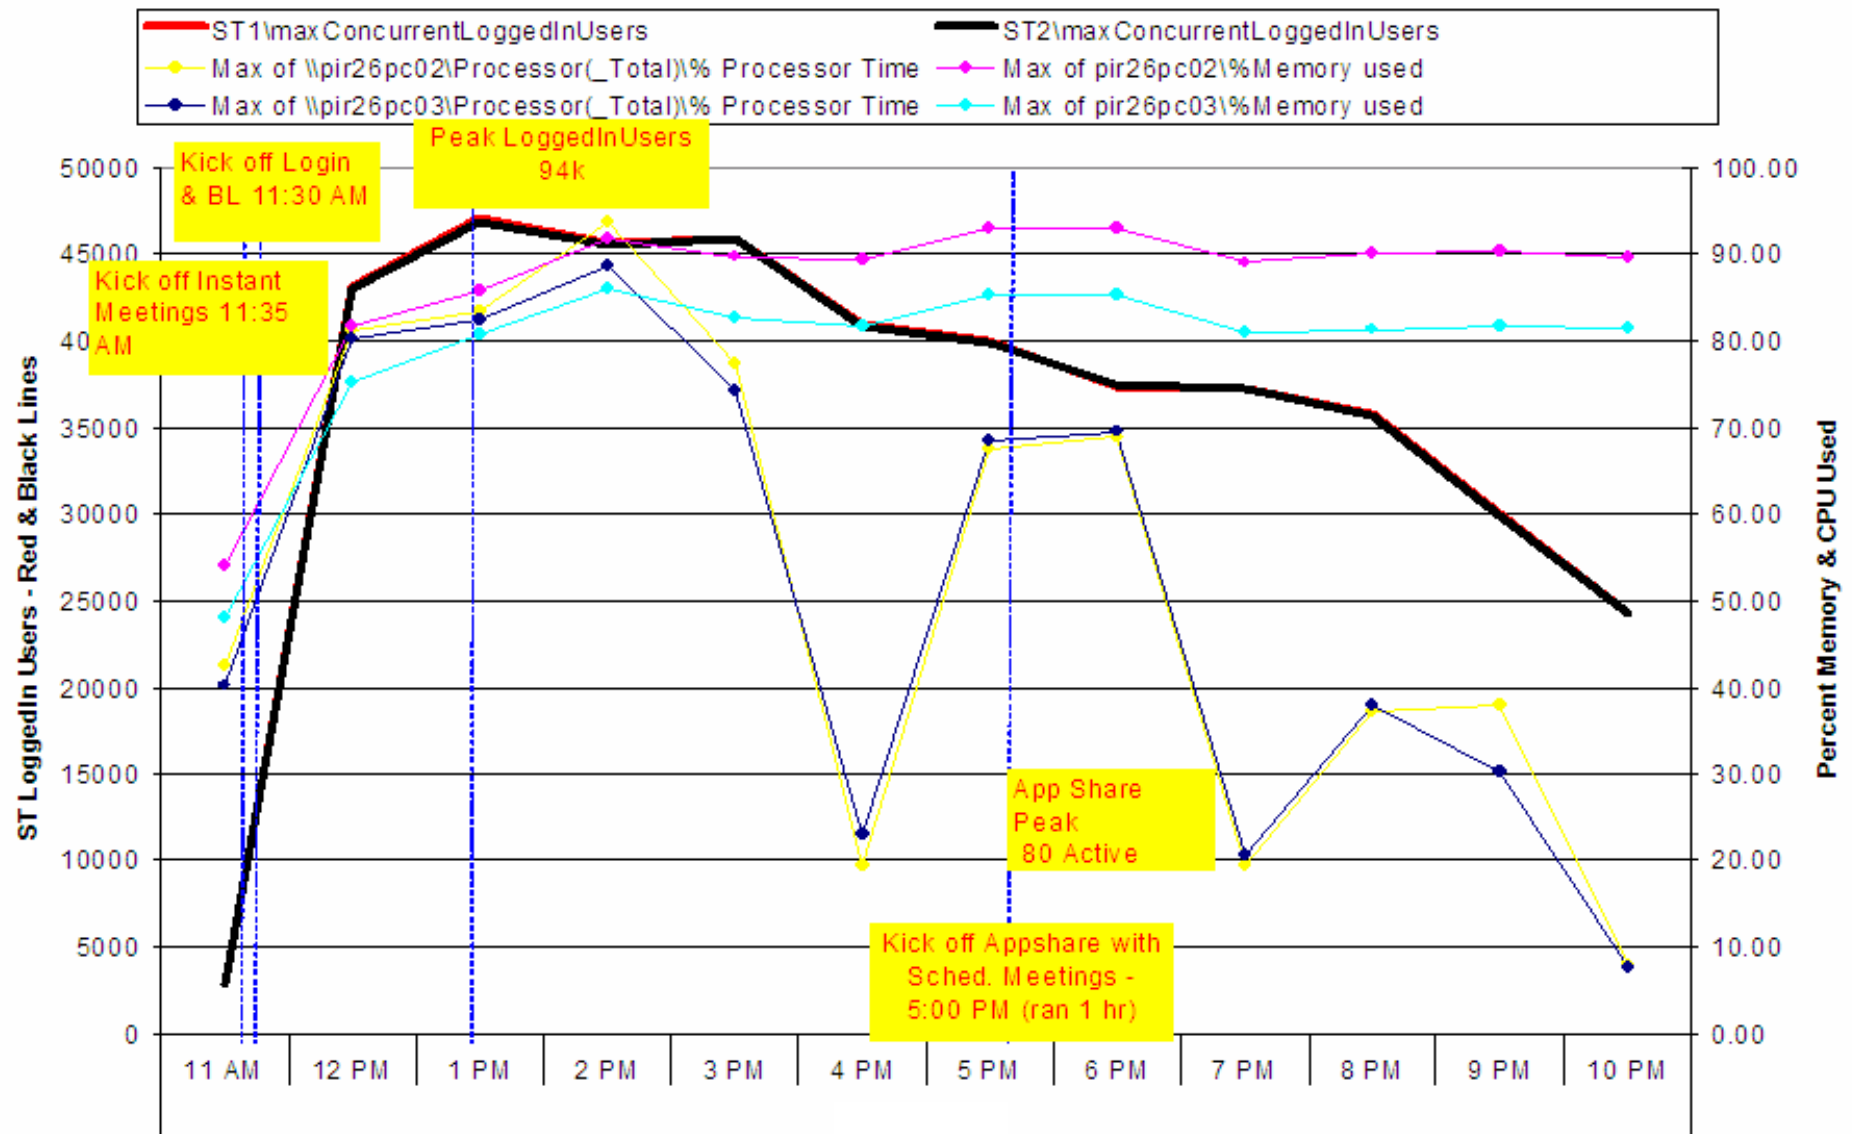
Testing Parameters

- Peak Logins: 94,000
- Peak Buddy List and Login: 68,000
- Buddy List Size: 15,367 entries
- N-Way Chat: 700
- Two Way Chat: 24,800
- Instant Meetings: 601
- Meetings with Application (Screen) Sharing: 782
- Users per meeting: 2 (always unique)

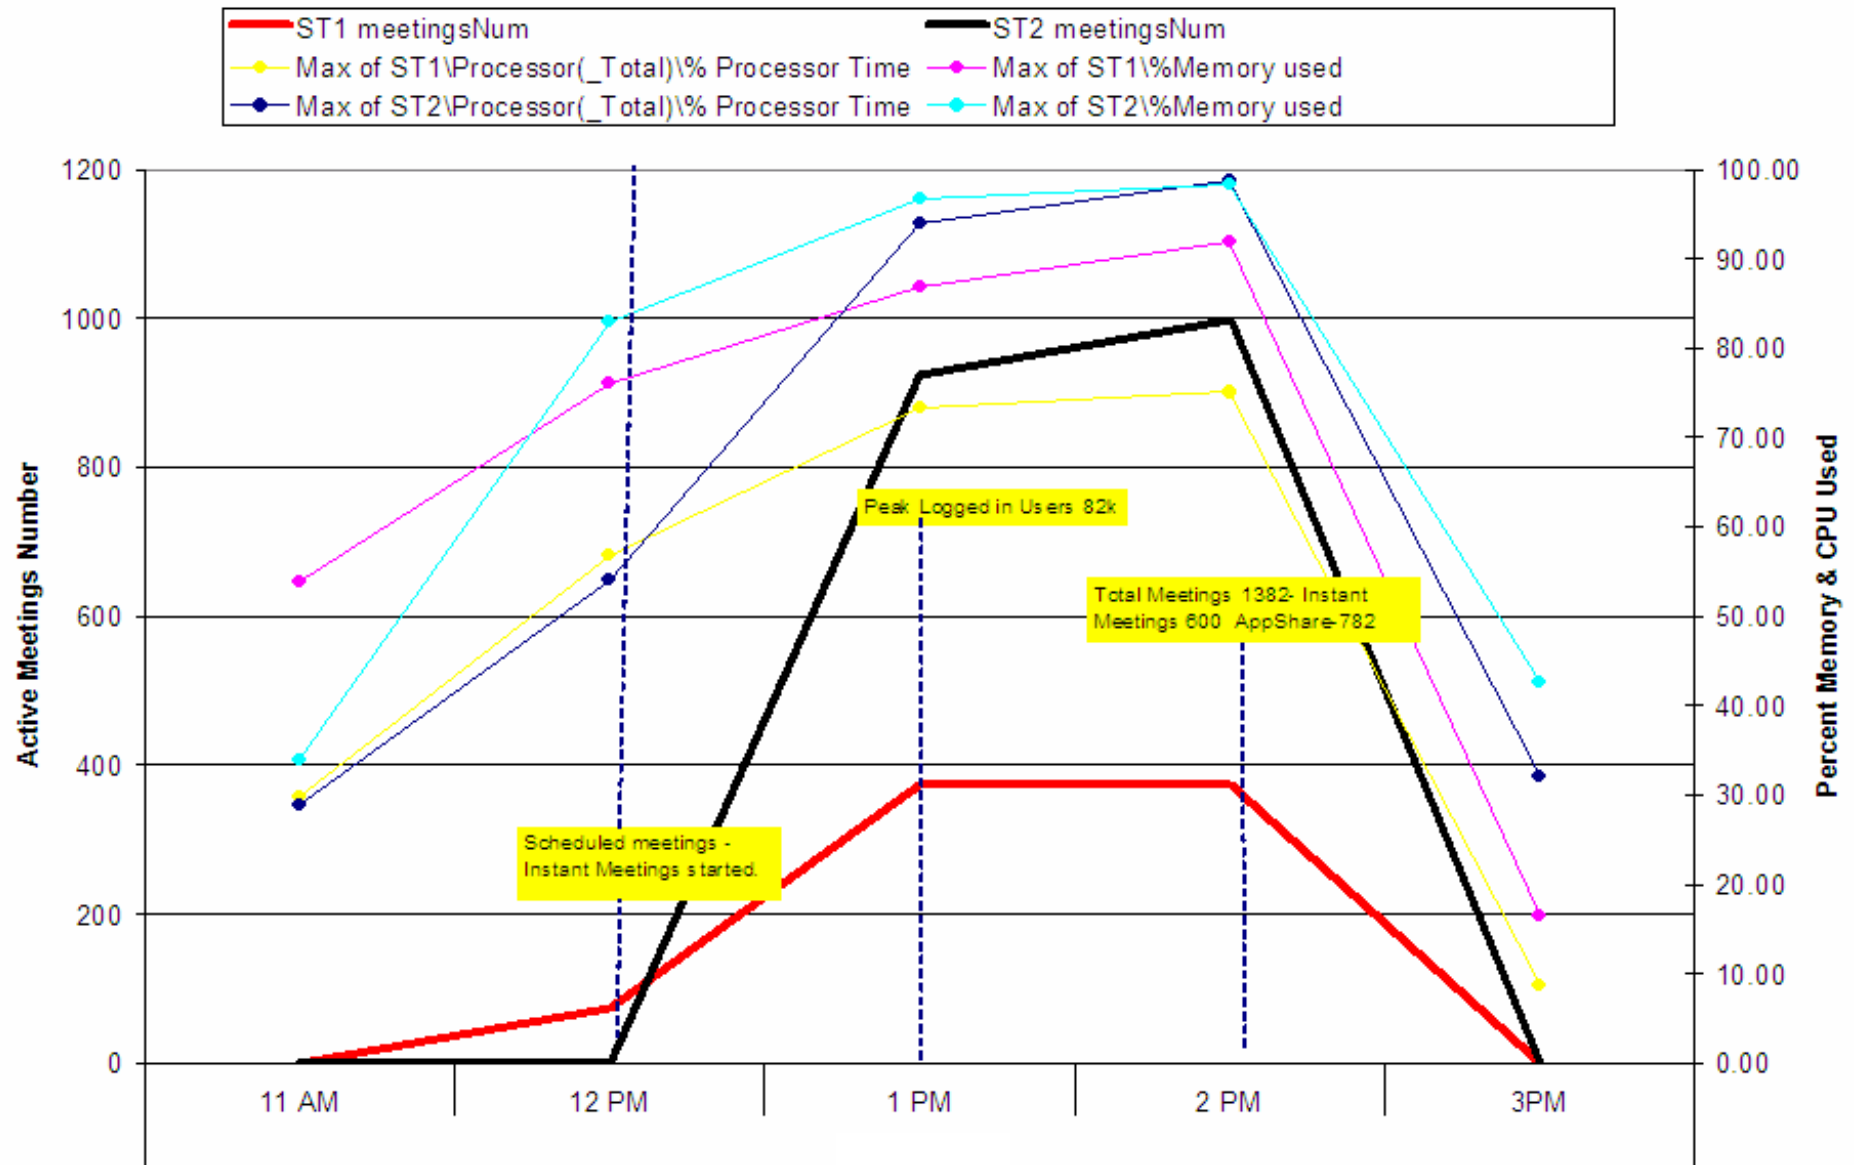
Community and MUX Server Testing

- Testing showed that with two MUX servers and two Community servers that:
 - ▶ A MUX server (4 way with 4GB of RAM) can handle at least 46,500 Presence Entities (users)
 - ▶ A Community Server cluster (2 servers) can handle 1,382 instant meetings and still handle 93,000 logons
- User Experience testing (i.e. manual testing) conducted whilst the automated testing , especially at peak level, was ongoing determined that performance for the user was within normal and acceptable limits and that the environment was acting as it would without the heightened load

Community Servers Memory / CPU



Instant Meetings / Screen (App) Sharing

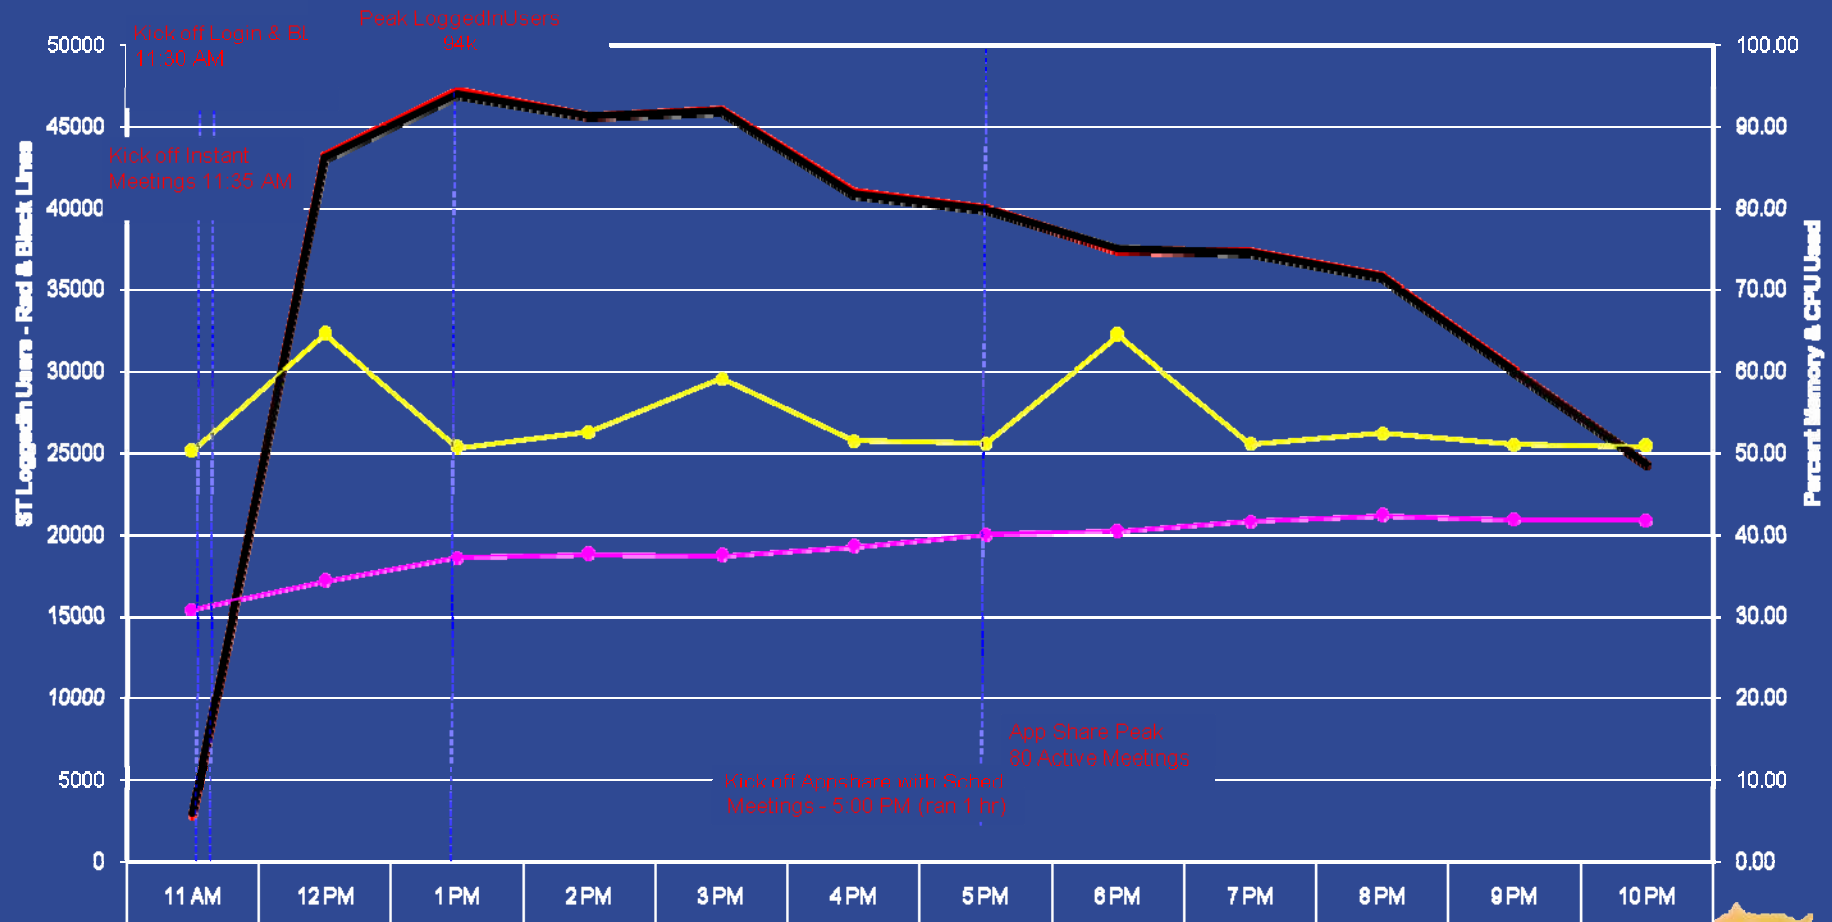


Lotus Sametime with LDAP

- Lotus Sametime comes with two directory choices:
 - ▶ Lotus Domino Directory
 - ▶ LDAP
- LDAP has proven to provide the most interoperability with extended applications and proven higher availability
- Many 3rd party products require LDAP directories to function.
 - ▶ With Lotus Sametime using the same LDAP directory they work out of the box.
 - ▶ If Lotus Sametime uses Lotus Domino Authentication you are required to do account mappings
- With proper monitoring and configuration you can setup an LDAP environment that is 100% available even with indexing, replication and crashes that generally plague systems

Lotus Sametime with LDAP Testing Results

ST1\maxConcurrentLoggedInUsers ST2\maxConcurrentLoggedInUsers Max of \LDAP1\Processor(_Total)\% Processor Time Max of \LDAP1\% Memory used

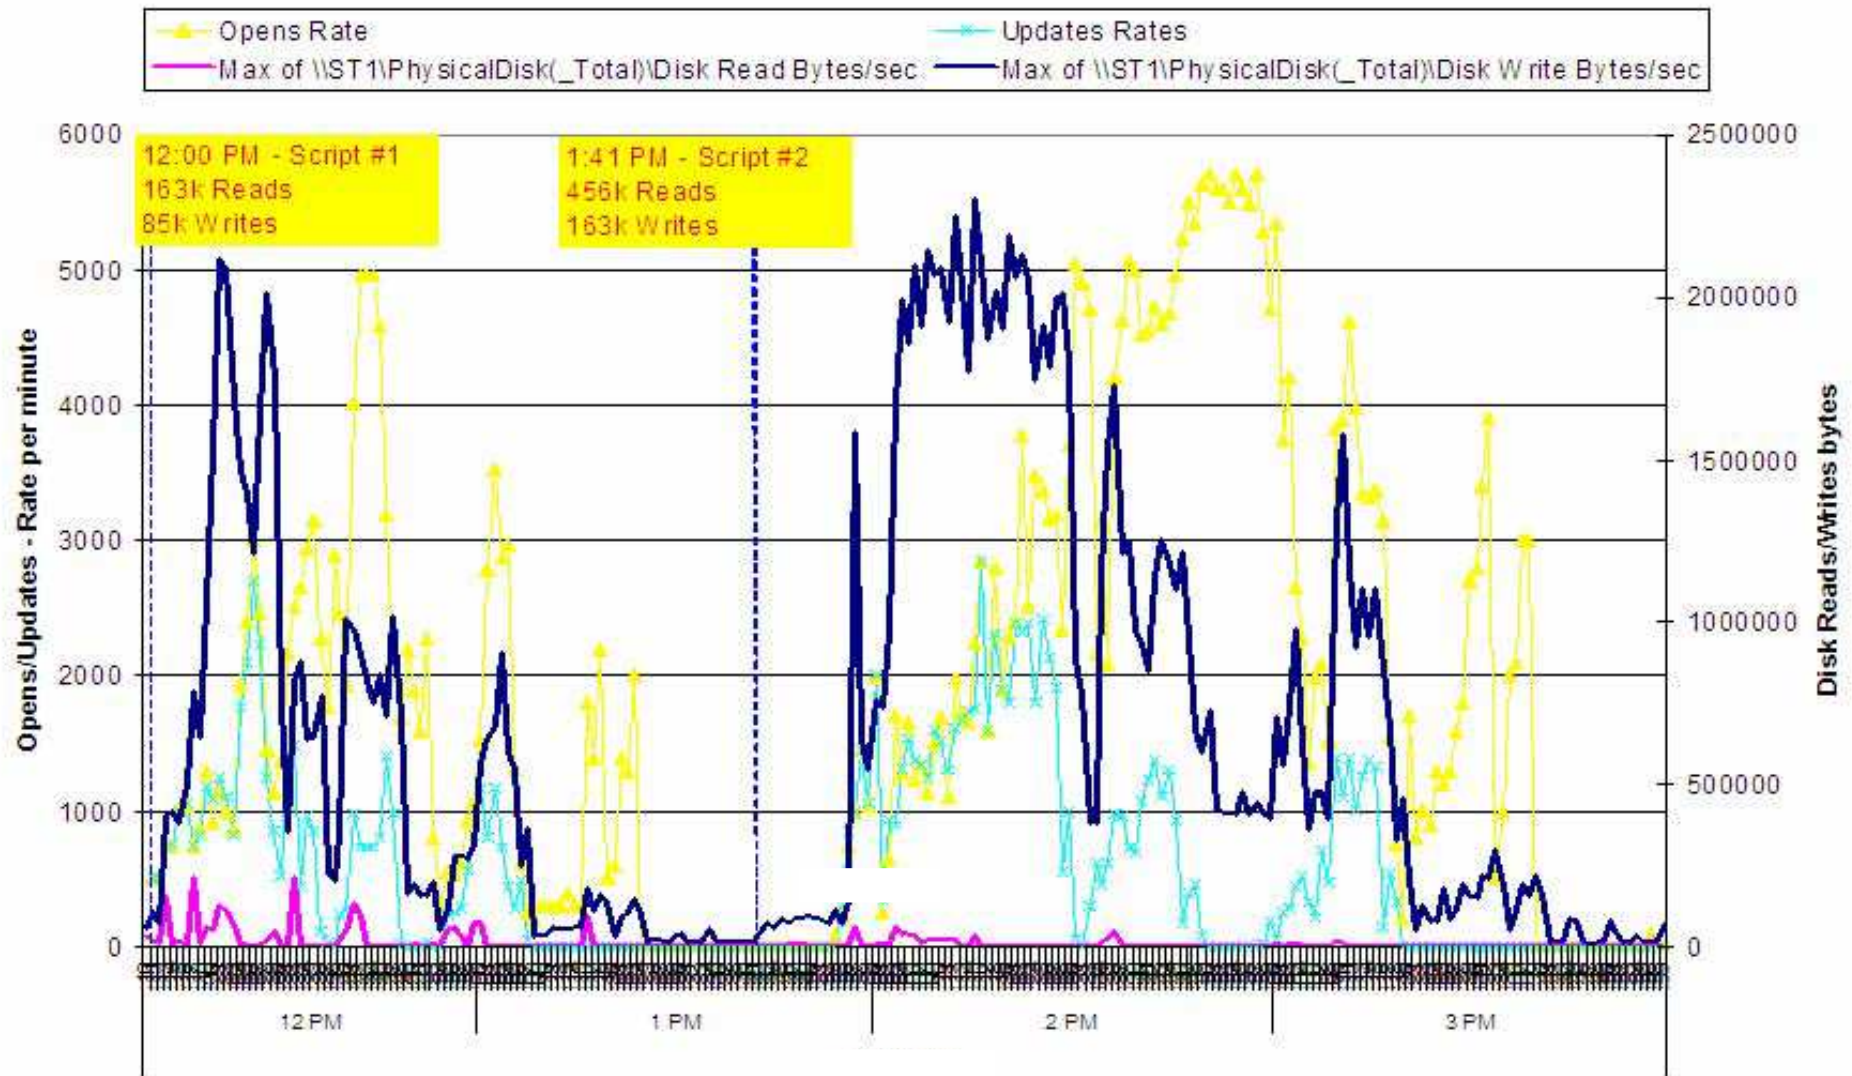


Lotus Sametime Buddy List Storage

- There have always been serious concerns about the servers ability to access, update and replicate thousands of documents in VUserInfo.nsf during peak usage times
- IBM has implemented a one off solution using DB2 to avoid these concerns based on testing with 4.6.5
- Customer wanted an “out of the box” easy to support solution
- Testing results showed that Lotus Domino 6.5.4 + could handle the load of STStorage
- Real world usage shows zero failures even with 85,000+ users connecting within a 5 min interval

ST VPuserInfo.nsf Testing Results

ST1 - vpuserinfo

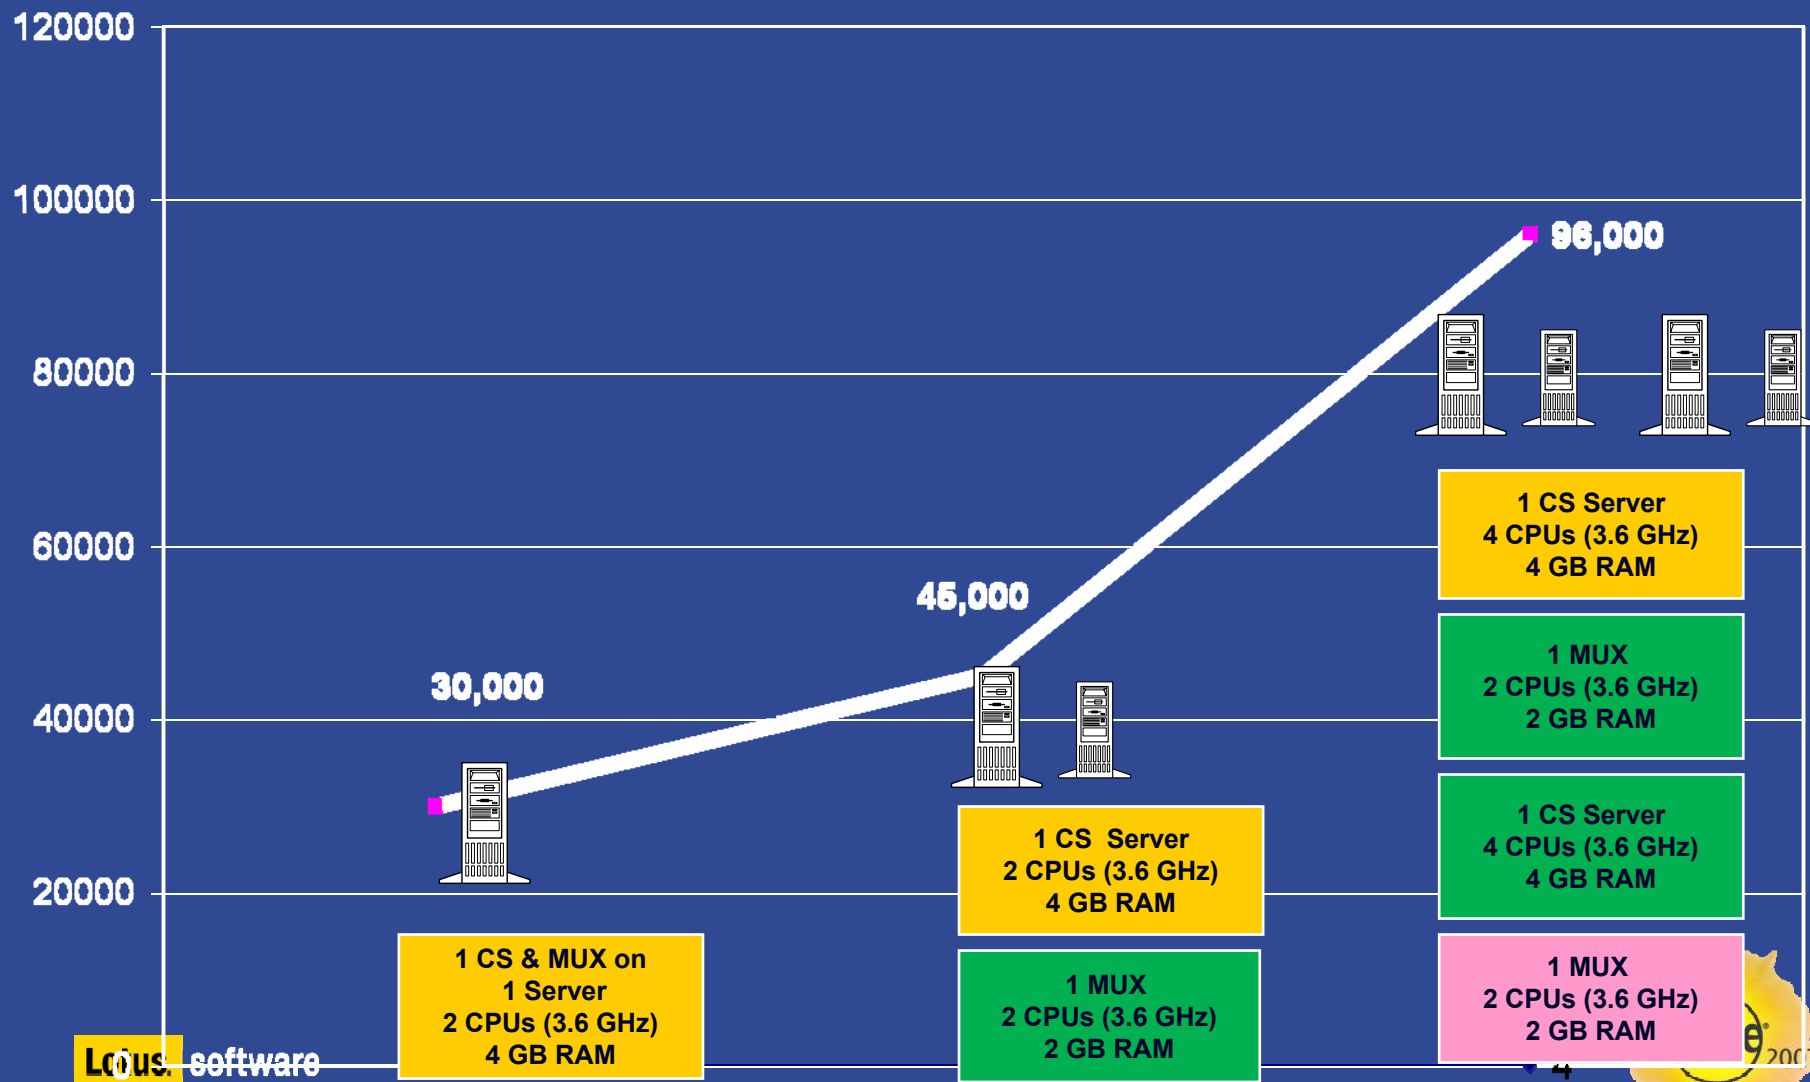


Bench Testing of Lotus Sametime 7.5

- Transaction Rates per user for 8 hour period
 - ▶ 2 logins
 - ▶ 16 userinfo requests for a random BuddyList users
 - ▶ 24 chats with a user on Buddylist – five 2-way messages exchanged with a partner
 - ▶ 18 status – change to random status and back to active status
 - ▶ 1 add a BuddyList member to existing BuddyList group
 - ▶ 1 resolve request
 - ▶ 1 awareness request
 - ▶ 2 policy retrivals on every login
 - ▶ 1 location set (initial login only)
 - ▶ 1 location BuddyList watch (initial login only)
 - ▶ 2 logouts

- Many of these numbers are based on real world situations at the Customer based on survey and monitoring data

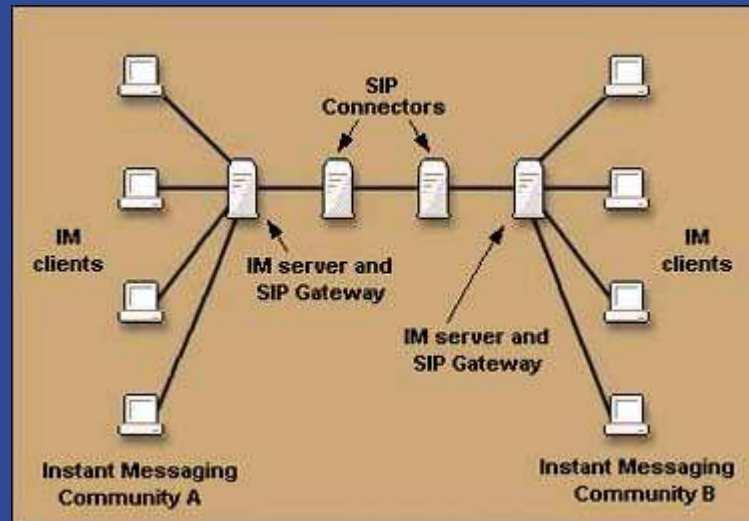
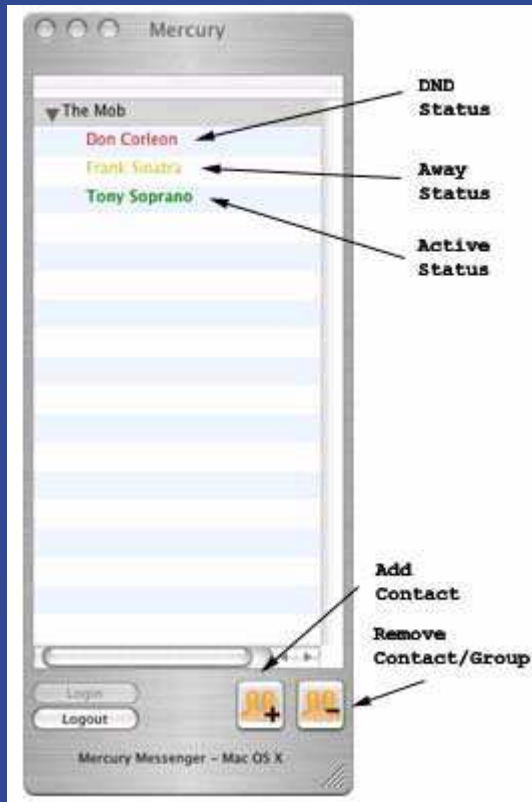
Testing of Lotus Sametime 7.5



Lotus software



Basic Architectural Principals





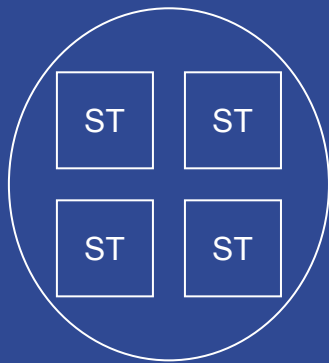
Deployment Architecture: Guidelines for Sizing

- Based on the latest testing on modern hardware as specified at a customer site on customer hardware:
 - ▶ MUX servers (i.e. stand-alone STMUX) can support at least 40,000 connections perhaps more
 - ▶ Conservatively, community servers can support at least 80,000 presence entities
 - ▶ A Community server can also support at least 400 instant meetings / application (screen) share sessions
 - ▶ A Community cluster with two stand-alone MUX servers can easily support 100,000 presence entities with more than 1,000 instant meetings
- However
 - ▶ Prior to this testing the rule was approximately 20,000 users per MUX server and 50,000 per community server with 100 instant meetings
 - ▶ Your mileage may vary
- Rate of use
 - ▶ For most customers, six months after deployment:
 - 95% of potential IM/awareness users logged in based on their time-slot
 - 3-5% of potential meeting users participating at any given time

Helpful Thoughts on Deployment

- Schedule updates to directory during off peak hours.
- Backup, purge and perform continual monitoring of stconf.nsf
 - ▶ Tools are available to help
- Buddy Lists
 - ▶ Use both cluster replication and scheduled replication of vuserinfo
 - ▶ Perform maintenance on vuserinfo by using IBM tools to validate names in buddy lists, etc.
- LDAP
 - ▶ Set cache for name lookup to 100MB
 - ▶ Implement a minimum of two LDAP servers per two community server cluster
 - ▶ Use at least three LDAP servers for a four community cluster

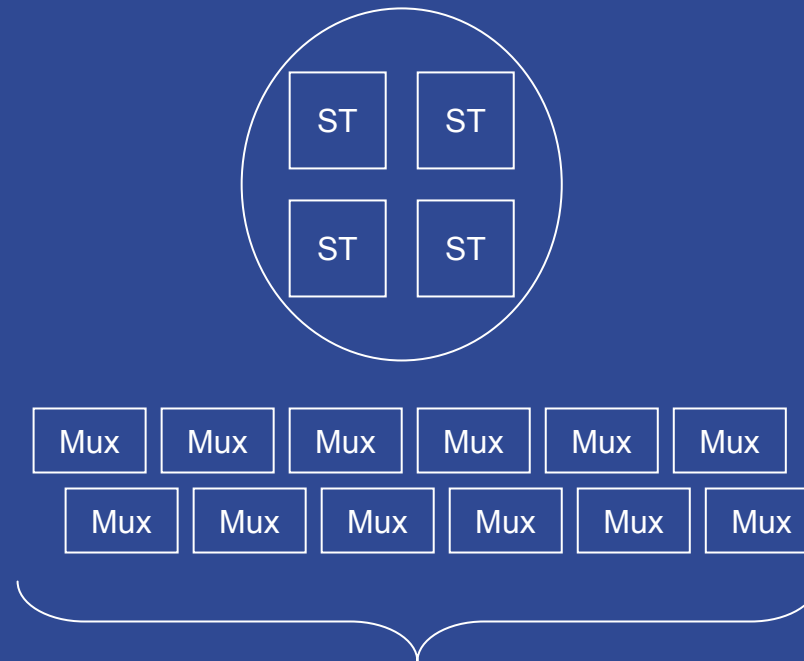
Real World Architectural Examples



IBM's Lotus Sametime Architecture

- ~ 300,000 Users
- 198,500 Peak Concurrent Persistence Entities
- Does not support Instant Meetings / Application Sharing

IBM's Lotus Sametime IM Architecture



330,000 users
in community

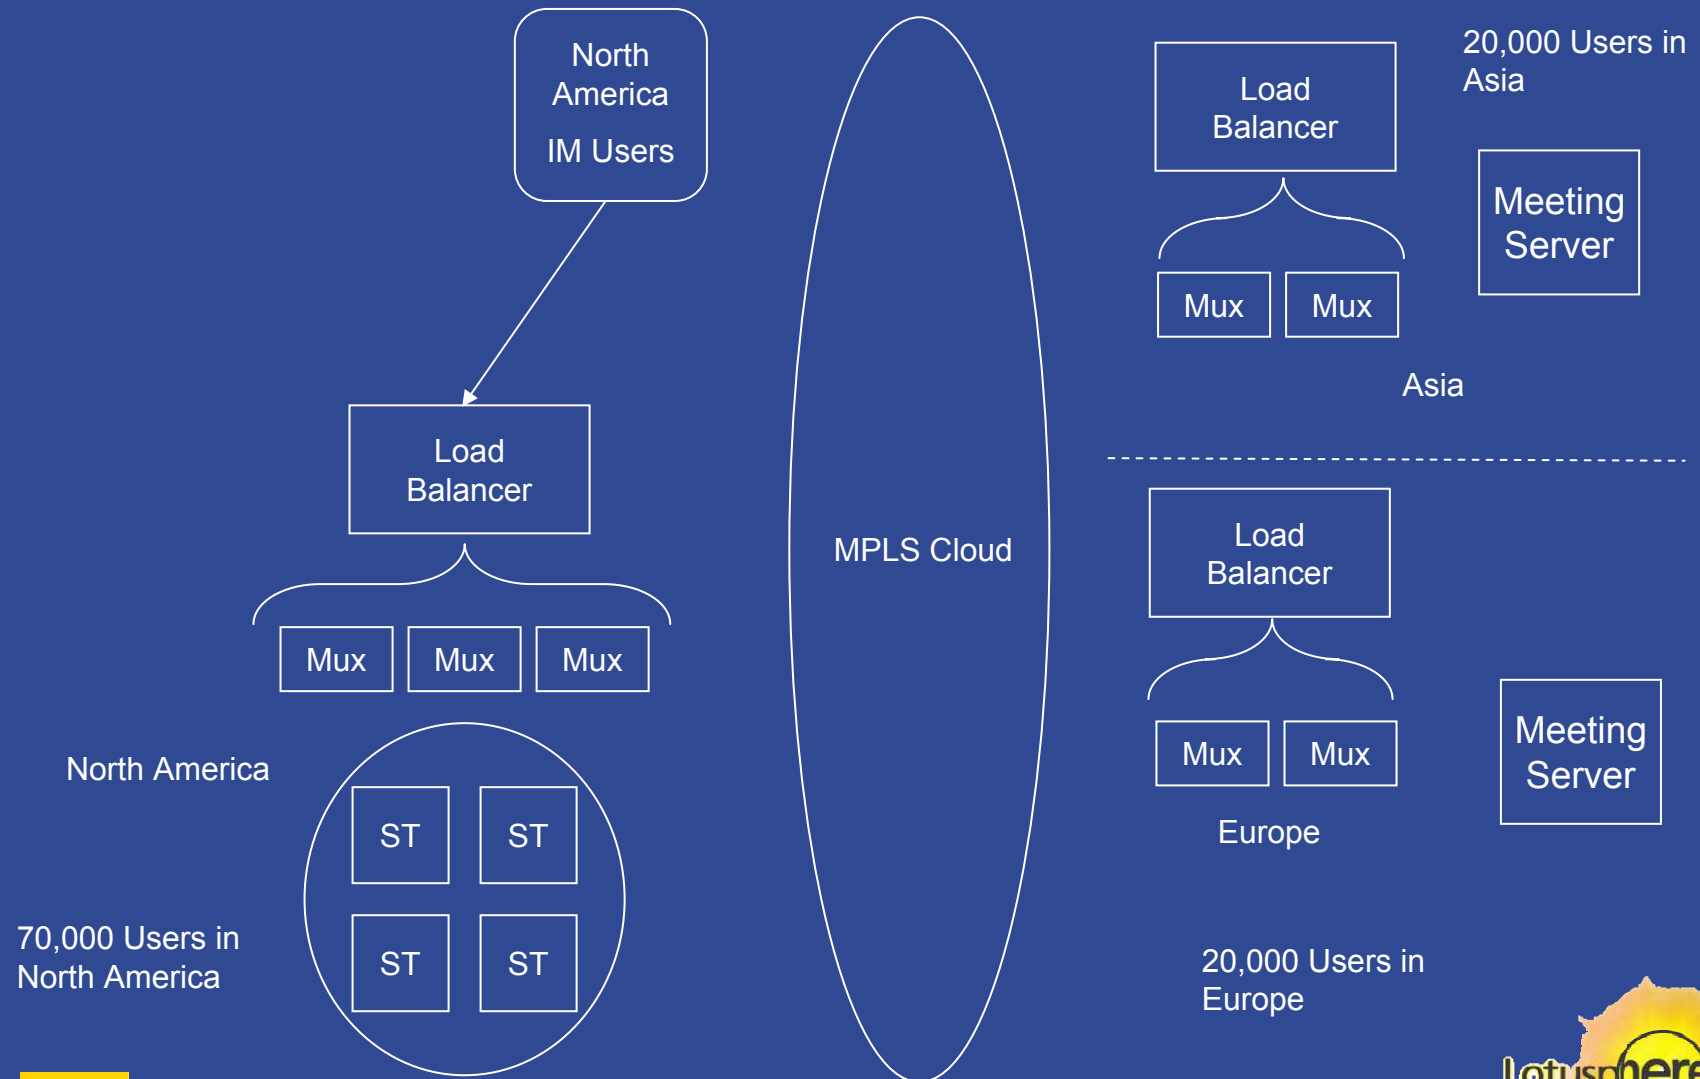
messaging.ibm.com
(round-robin DNS)

Peak concurrency
198,500 users

Large Hybrid Architecture

- ST 7.0
- Sun One LDAP
- 105,000 Users
- 51,000 Peak Concurrent Persistence Entities

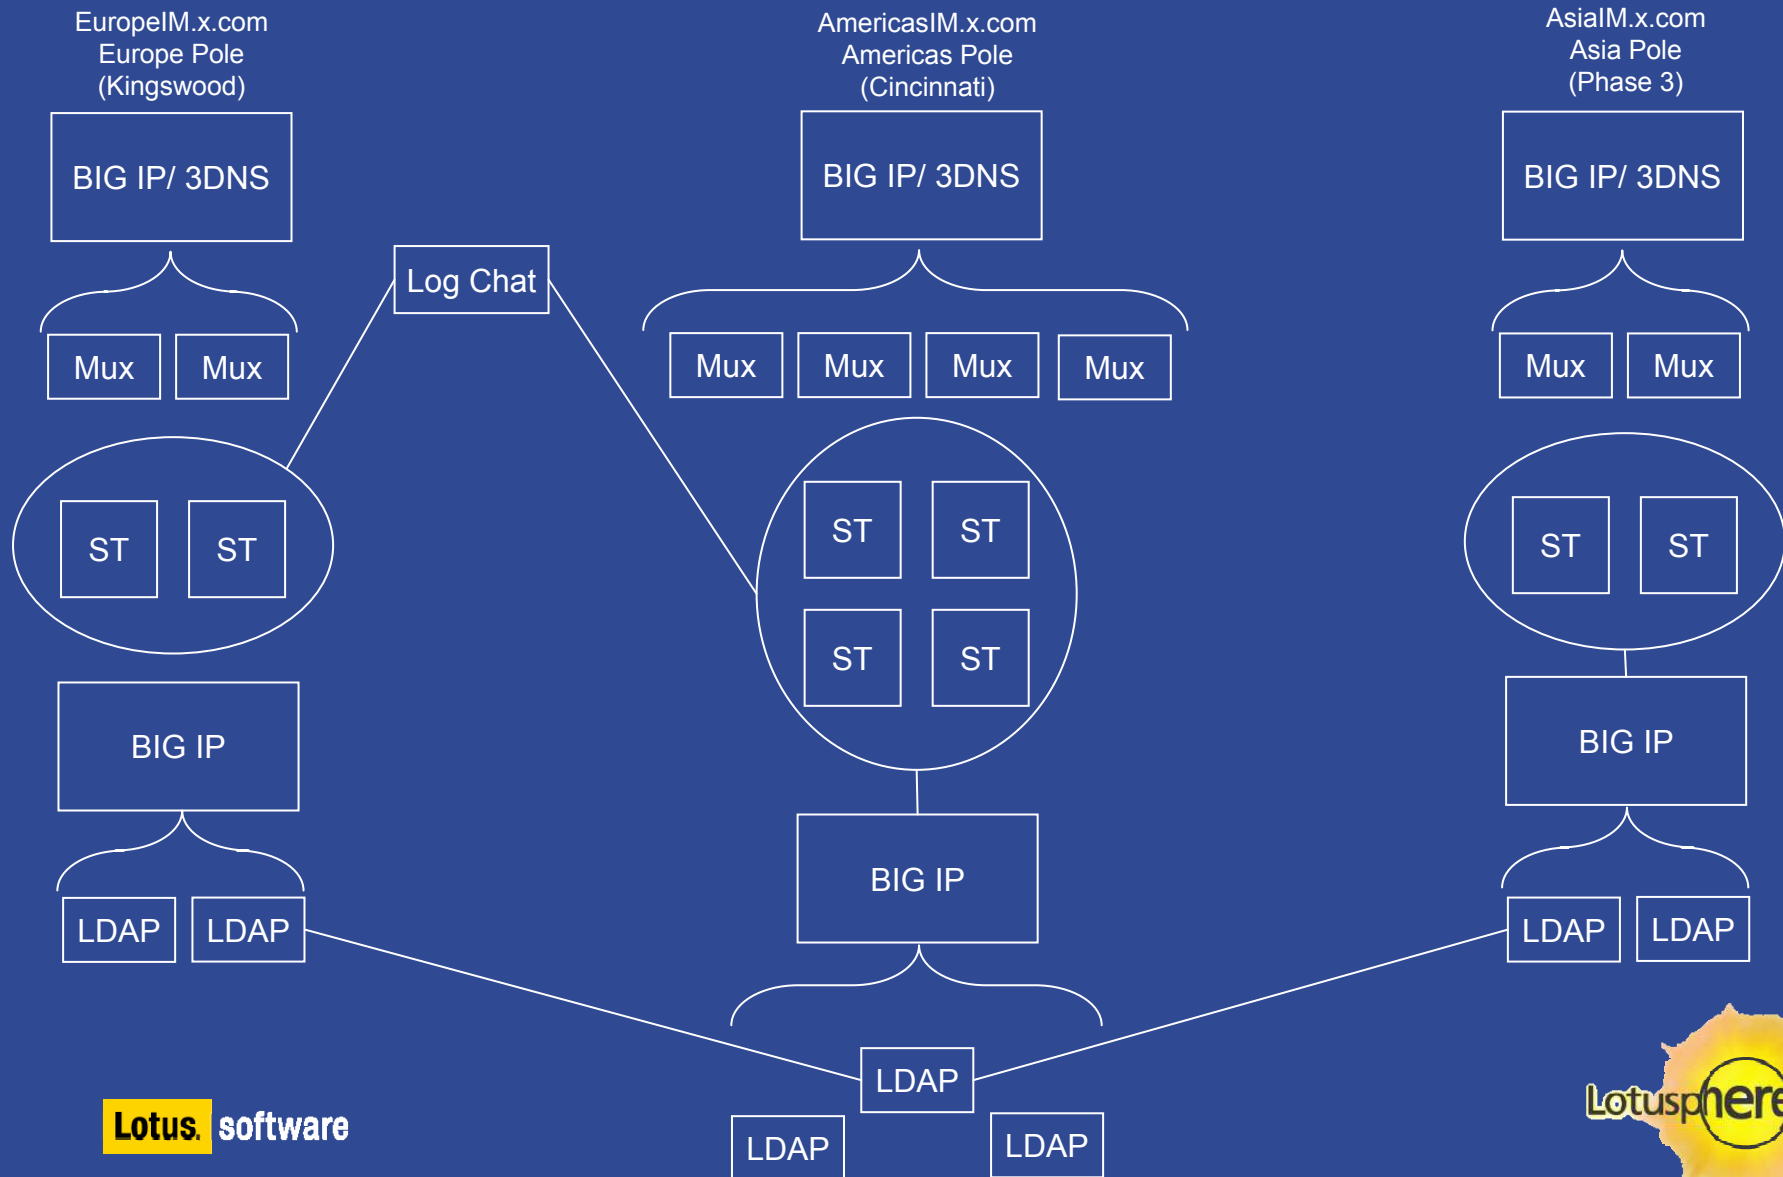
Large Deployment Hybrid Architecture



Customer's Lotus Sametime Architecture

- Lotus Sametime 6.5.1 FP1
- 360,000 Users
- 201,000 Peak Concurrent Persistence Entities
- Supports Instant Meetings and all Lotus Sametime functions
- Scheduled Web Conferences occur on a separate infrastructure that uses 32 stand-alone room servers

Customer's Lotus Sametime IM Architecture



Advantages to Customer's Architecture

- Provides IM and Instant Meeting functionality to their end users
- Clusters are designed to operate at 50% hardware availability with 100% functionality
- Architecture takes advantage of the network between poles. Eight community servers connections across the WAN is much better than 35,000 end user connections!
- Designed with 7.5 in mind. Pole design is ready for “Click to Talk”, secure file transfer and “Click to Call” functions by keeping as much data as possible local

Customer's Results

- Europe availability increased from 92% to 100%! There has not been a single trouble ticket in Europe since May of 2006!
- Americas availability has been 99.996%. Several configuration problems with BigIP recently addressed and availability is climbing
- No longer called Sometimes by many end users. Usage has increased by ~ 40,000 users since the deployment
- Failures in the infrastructure have proven the architecture many times over. A BigIP failure can send all 85,000 Americas users back to the Community servers all at the same time. Customer can get all these users reconnected and buddy lists resolved in less than 5 minutes!
- 3 LDAP servers handle 180,000 requests a min during peak times. Have serviced Customer's entire collaboration environment during an emergency

IBM

IBM Corporation 2007. All Rights Reserved.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

IBM, the IBM logo, Lotus, Lotus Notes, Notes, Domino, Sametime, WebSphere, Workplace and Lotusphere are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Lotus software

Other company, product, or service names may be trademarks or service marks of others.

