

What you need to know for the final

Final Exam Topics (Very approximate)

- 20%: Crypto/authentication
- 10%: Firewalls & VPNs
- 10%: Network concepts
- 10%: Distributed file systems
- 8%: Distributed shared memory
- 8%: Clustering
 - + RPC, clocks, mutex, deadlocks, concurrency, distributed transactions, fault tolerance

Distributed Systems
Spring 2009
Paul Krzyzanowski

Systems

- Bus versus switched interconnect
- Snoopy cache

Distributed Systems
Spring 2009
Paul Krzyzanowski

Networking

- Connectionless vs. connection-oriented
 - Datagram, virtual circuit
- Broadband versus Baseband
- Network layer versus Transport layer
- Internet Protocol: UDP/IP, TCP/IP Sockets
- Protocol Control Block

Distributed Systems
Spring 2009
Paul Krzyzanowski

Remote Procedure Calls

- Language-level construct
- Functional RPCs:
 - Sun, DCE
- Object-oriented RPCs:
 - Microsoft DCOM/ORPC, Java RMI, CORBA
- XML-based RPCs and Web services:
 - SOAP, XML RPC
 - AJAX (JavaScript + XMLHttpRequest), REST

Distributed Systems
Spring 2009
Paul Krzyzanowski

Distributed File Systems

- Stateful versus stateless design
- Upload/download versus RPC model
- NFS
- AFS
 - whole-file upload
 - Session versus sequential semantics
- Coda
 - Read/write volume replication
 - Disconnected operation
- DFS
 - Consistency tokens
- SMB/CIFS
 - Strong consistency
 - Oplocks

Distributed Systems
Spring 2009
Paul Krzyzanowski

Clocks

- Logical clocks
 - Event ordering
 - Lamport timestamps
 - Vector clocks
- Physical clocks
 - Cristian's algorithm
 - Berkeley synchronization
 - NTP/SNTP: synchronization subnet

Distributed Systems
Spring 2009
Paul Krzyzanowski

Mutual Exclusion

- Centralized
- Ricart & Agrawala
- Lamport
- Token Ring

Distributed Systems
Spring 2009
Paul Krzyzanowski

Distributed Shared Memory

- Implementation
 - Page fault on MMU
 - Page residence maintained by directory
- Sequential consistency
- Weak consistency models
 - Barrier
 - Release
 - Entry

Distributed Systems
Spring 2009
Paul Krzyzanowski

Fault Tolerance

- Redundancy
 - Physical (e.g., TMR)
 - Information (e.g., Hamming codes, RAID-4/5)
 - Temporal (retransmission)
- Byzantine faults versus fail-silent faults
- Two-army problem

Distributed Systems
Spring 2009
Paul Krzyzanowski

Cryptography

- Symmetric versus public key
- Key length & brute-force attacks
- Key exchange
 - Third-party arbiter, Diffie-Hellman, Public Keys
- Hash functions
- Digital signature (encrypted hash)

Distributed Systems
Spring 2009
Paul Krzyzanowski

Authentication

- Nonce-based
- Authentication + key exchange
- Certificate: signed public key
- Kerberos
- Biometrics
 - Statistical threshold
- OpenID

Distributed Systems
Spring 2009
Paul Krzyzanowski

Security

- Problems
 - Buffer overflow
 - Denial of Service (DoS)
 - SYN flooding
 - Worms, viruses, key loggers, rootkits
- Approaches
 - Authentication, Authorization, Accounting, Auditing
 - Identification vs. Authorization
 - Code signing

Distributed Systems
Spring 2009
Paul Krzyzanowski

Firewalls

- Screening router (packet filter)
- Application proxy
- DMZ (screened subnet)
- Bastion hosts
- VPN: tunneling (+ encryption + signatures)

Distributed Systems
Spring 2009
Paul Krzyzanowski

Clusters

- Types:
 - HPC
 - Batch processing
 - Load balancing
 - High-availability
- Shared disk versus shared-nothing
 - DLM (distributed lock manager)
- Warm vs. cold failover
- Multi-directional & Cascading failover
- Load balancing

Distributed Systems
Spring 2009
Paul Krzyzanowski

Virtualization

- Storage virtualization
- Virtual machines
 - VMM (hypervisor)
 - Privileged vs. unprivileged instructions

Distributed Systems
Spring 2009
Paul Krzyzanowski

The End.

Distributed Systems
Spring 2009
Paul Krzyzanowski