


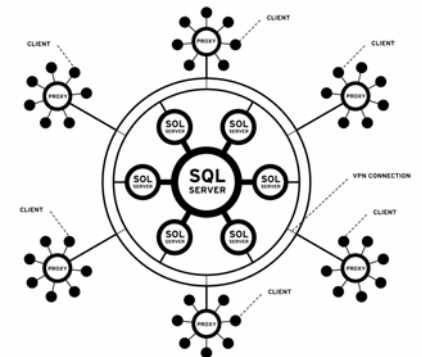


# Scaling EVE Online

## Under the Hood of the Network Layer




Davíð H. Brandt  
Senior Network Programmer  
EVE Online  
CCP Games








## The Vision

- A social & gaming theory
- MMORPG vs UMMORPG
- 20.000+ concurrent users per cluster
- 100.000+ registered users per cluster
- Avoid sharding at all costs





## Sharding as a Scalability Model




- A subset of the MMORPG world and user base
- Single logical server + database
- Typically  $< 3.000$  users per shard
- To add more users, add more shards
- Reduces the complexity of distributed programming
- Simplifies content reuse



## Single-Shard Scalability

- Entire user base shares hardware
- Single physical location on the Internet
- Single logical database
- Lowest Common Denominator users
- User state handover & sharing
- Increased programming complexity
- User-perceived performance
- TCO per concurrent user
- Hot-spots & Load Balancing
- The “O(Server Count)” problem
- The “Server Maintenance” problem





## August 2002: The Challenge

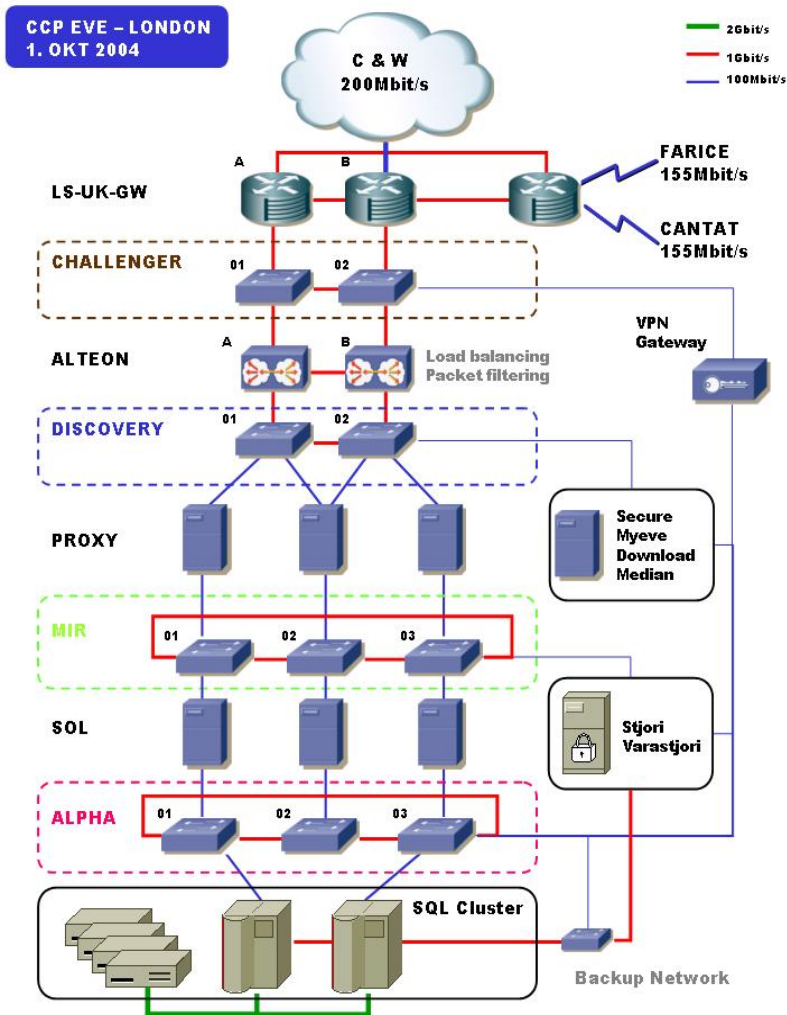


- The first “UMMORPG”
- Scalability as an afterthought
- Legacy services, session and security
- Lowest Common Denominator
- Lack of time & manpower
- Kamikaze Project: MachoNet

*“It's 106 miles to Chicago, we've got a full tank of gas, half a pack of cigarettes, it's dark and we're wearing sunglasses.”*



## The Solution: Hardware



- 400 GHz CPU / 200 Gb RAM
- 2 Routers (CISCO Alteon)
- 14 Proxy servers (IBM Blade)
- 55 Sol servers (IBM x335)
- 2 DB servers (clustered, IBM Brick x445)
- FastT600 Fiber, 56 x FC 15k disks, DS4300 + 3\*EXP700
- Windows 2000, MS SQL Server



## The Solution: Software

- Full-fledged proprietary OORPC
- Distributed Session Management
- Multicasting & Complex Addressing
- Predictive Load Balancing
- Scalable Services:
  - Solar Systems
  - Corps & Alliances
  - Markets
  - Chat






B.O.F.H.



24/7-users



Corporate Firewall

## • Lowest Common Denominator

- Firewalls, Lossy Networks

## • Diverse protocol requirements

- “Excel 3D” vs Space Game

## • MMORPG vs FPS Lag

## • Network vs User-Perceived Lag

## • UDP related support costs



N.A.T.

¿En Un Atasco?



Congested Networks  
(packet loss)



P.F. Admin




Users in Network Hell



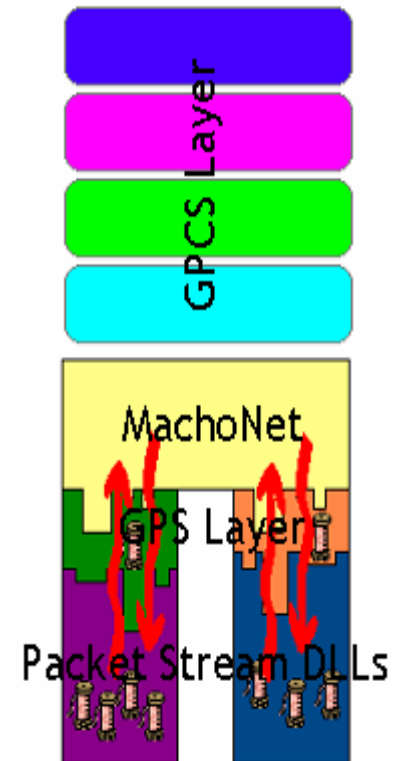
Personal Firewall



# MachoNet: The Conventional Parts

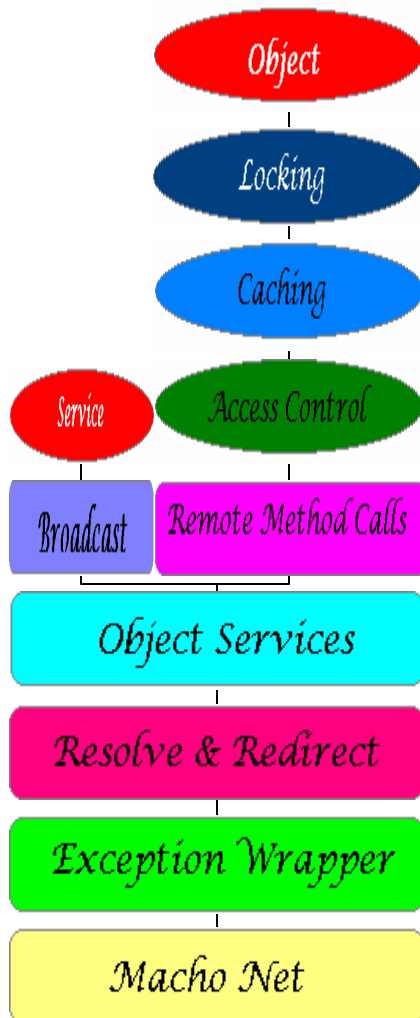


- 800 remote-methods
- 71 stateless/semi-stateless services
- 28 stateful services (“bound objects”)
- 80%-99% “non-game”
- All the usual OORPC features
  - Classical RMI syntax
  - Remote garbage collection
  - PBR / PBV
  - Access control (role/method)
  - Remote Exceptions
  - Huge method
  - Thread pooling
  - Object Lookup
  - Statistics & Logging





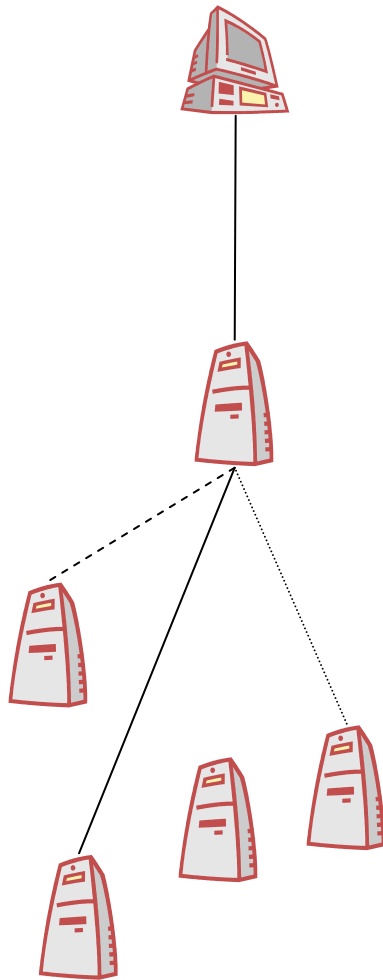
## MachoNet: The Unconventional Parts I



- IDL free
- Method Call Caching
- Blocking vs Non-Blocking
- Implicit locking
- Preemptive vs Non-Preemptive
- Game Session based security
- Game Session-change awareness




# MachoNet: Distributed Session Management



- Client session is mastered on Proxies
- Contents:
  - User, character & security roles
  - Groupings (alliance, corporation, gang)
  - Location (region, constellation, solar system, station, ship)
  - Object References
- JIT Session State Propagation
  - Session Fragmentation
- Availability:
  - RMI
  - Worker Threads
- Applications:
  - Security
  - Multicasting
  - Game logic



## MachoNet: The Unconventional Parts II

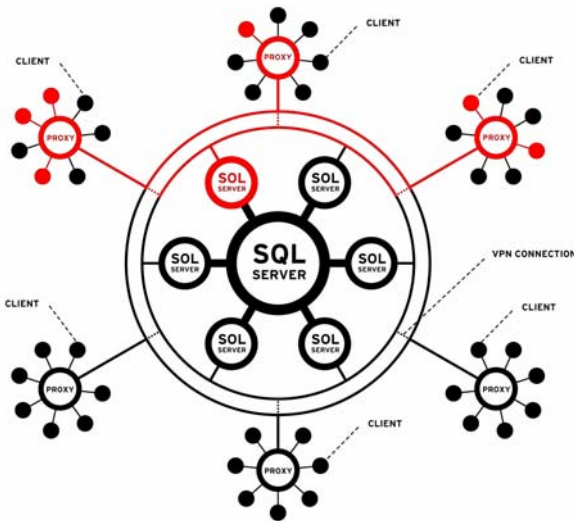


- Multicasting
- Ordered & Unordered
- Reliable Delivery
- Multicast bundling
- Complex addressing
  - eventName & (nodeID | proxyID | [session (attribute, value)]+)



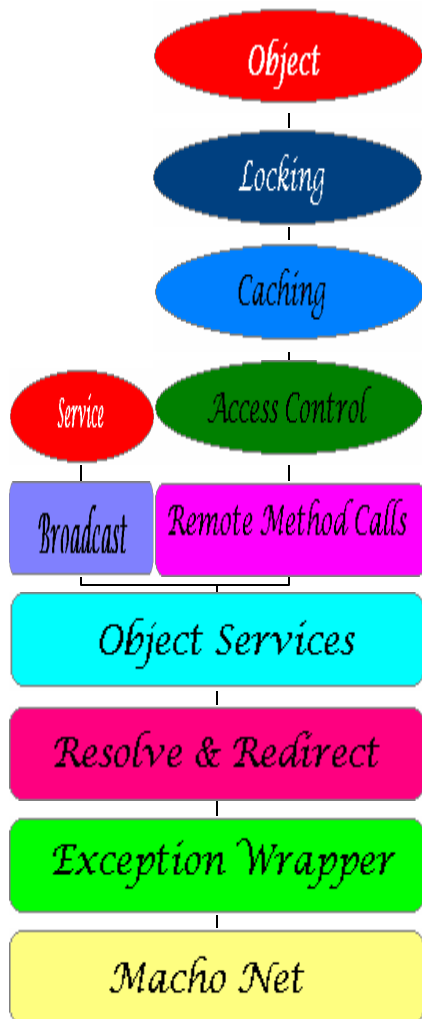
## MachoNet: Multicasting via Proxies

- Why not true multicasting?
- The “O(Proxy Count)” problem
- Multicasting by session info
  - On “sol” node
  - On “corp” node
  - On “wrong” node
- Sequenced Multicasting






## MachONet: The Unconventional Parts III



- Load Balancing features
  - Any Node
  - Session Derived Node
  - Mapped/Bound Objects
- Complex addressing
  - Mapping Address | nodeID+Instance
  - serviceID & ( nodeID | any node | session derived node | all nodes)
- Method bundling
  - QueuedCast
  - Batched Remote Calls / Scatter-Gather



# MachoNet: Predictive Load Balancing



- Load estimation
- Cluster segmenting
  - Markets
  - Corp / Alliance
  - Solar Systems by Constellation
  - Stations by Solar System
- CPU per node
- RAM per host
- The “Yulai/Rens” Problem
- The “Fleet Battle” Problem
  - Dynamic Load Balancing





## Service Scalability: Corps & Alliances



- Mapping by CorpID / AllianceID
- The “Session Convergence” Problem
- Impact on stability
- Resource utilization vs Stability
- Impact on interactions & gameplay



## Service Scalability: The Market



- Mapping by RegionID
- Solving “Session Convergence”
  - The Market Proxy
- Caching
- Impact on notifications
- Vertical scalability issues
- Impact on supply & demand



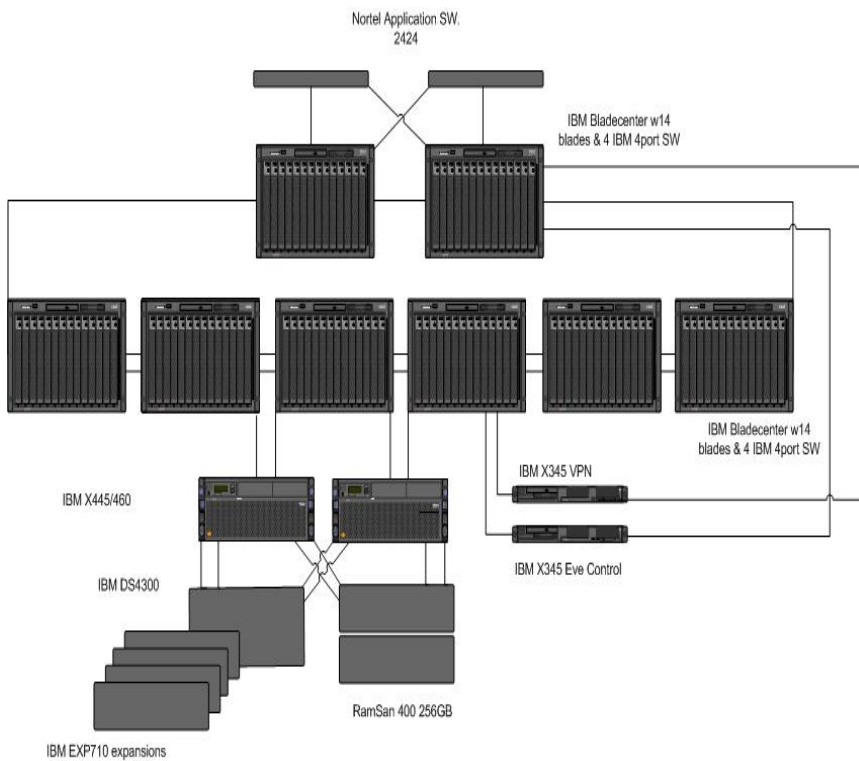
## Service Scalability: The Chat System



- LSC vs Predecessor
- Chat with 500-1000 participants
- Solving “Session Convergence”
  - Distributed Stateless Chat
- The “Packet Bus”
- Member Lists
  - The “O(Member Count)” Problem
  - Immediate Mode & Batched Method Calls
  - Delayed Mode & Scatter-Gather bookkeeping
- The “Local Chat” Issue



# The Future of EVE Online



- Content & Graphics
- Hardware (RAMSAN)
- Sharding
  - China: ~30k
  - London: ~25k
- Algorithmic Improvements
- Dynamic Load Balancing
- 64-bit Windows
- EVE Mobile



# EVE

WWW.EVE-ONLINE.COM