Game Developers Conference™ Europe 2011
August 15-17, 2011 | Cologne, Germany
www.GDCEurope.com

Inside Tibia
The Technical Infrastructure of an MMORPG

Matthias Rudy
Head Programmer, CipSoft GmbH

Tibia

- 2D fantasy MMORPG for PC
- online since 7 January 1997
- commercial since 5 November 2001
- free to play
- optional subscription (7.50 Euro for 30 days)
- some paid extra services
  - world transfer, name change, but no ingame items
- two clients
  - stand-alone client for Windows and Linux
  - Flash based client for browsers since June 2011

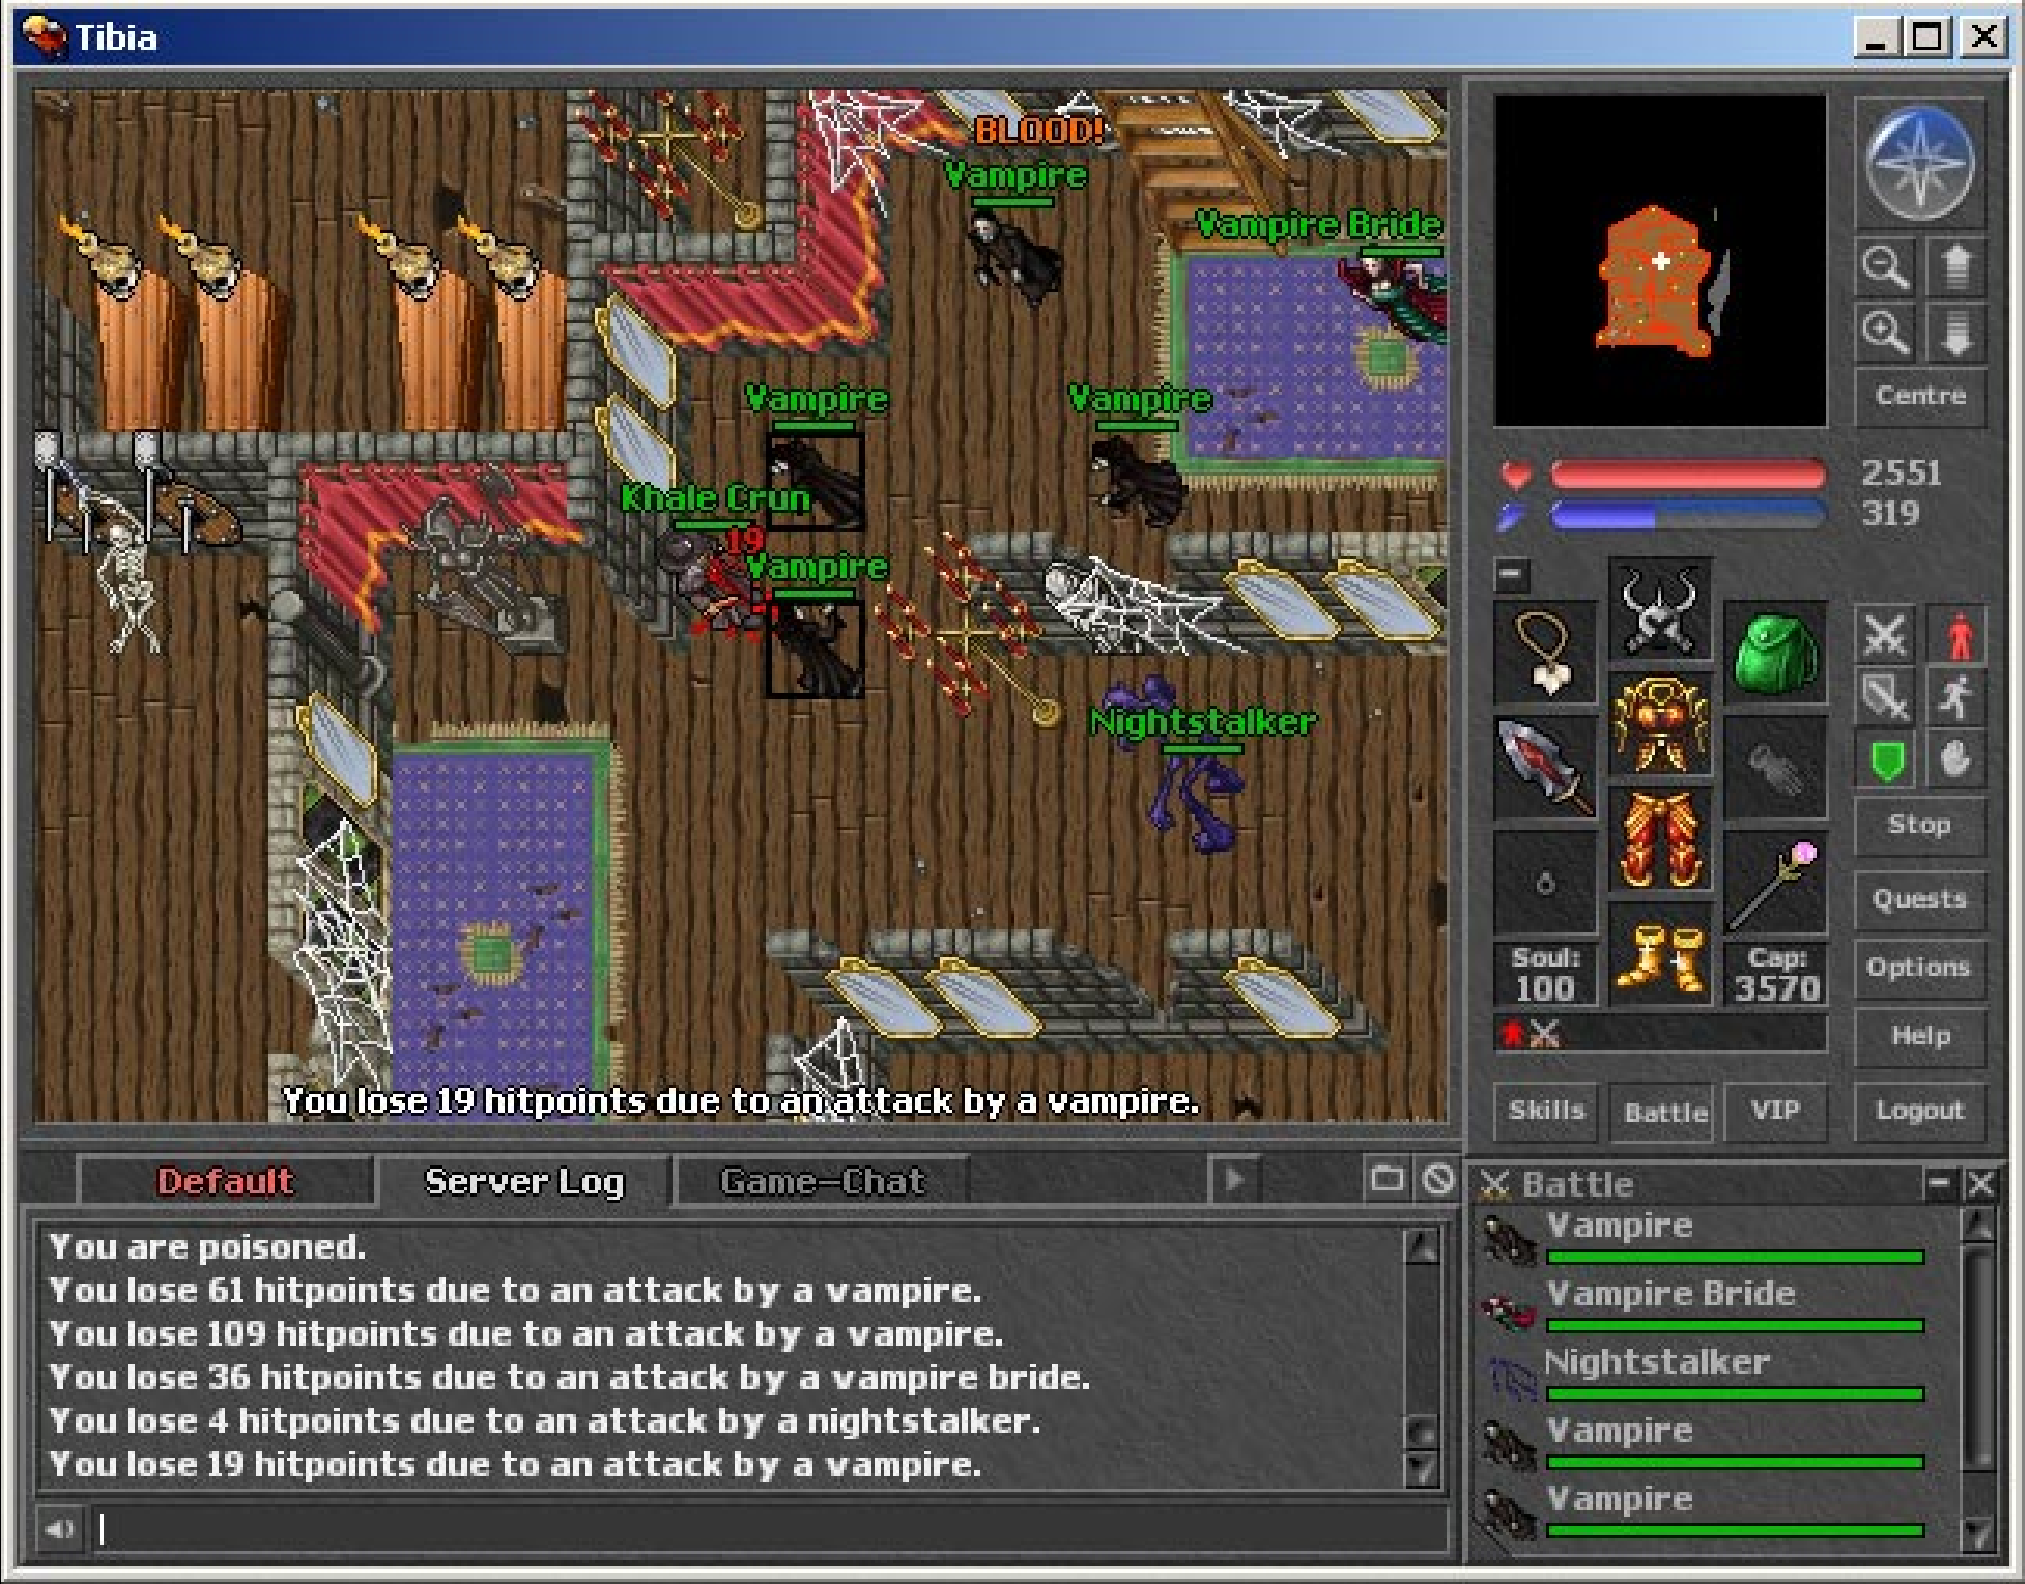
Tibia
BLOOD!
Vampire
Vampire Bride
Vampire
Vampire
Khale Crun
19
Vampire
Nightstalker
You lose 19 hitpoints due to an attack by a vampire.
Centre
2551
319
Soul:
100
Cap:
3570
Stop
Quests
Options
Help
Skills
Battle
VIP
Logout
Default
Server Log
Game-Chat
You are poisoned.
You lose 61 hitpoints due to an attack by a vampire.
You lose 109 hitpoints due to an attack by a vampire.
You lose 36 hitpoints due to an attack by a vampire bride.
You lose 4 hitpoints due to an attack by a nightstalker.
You lose 19 hitpoints due to an attack by a vampire.
Battle
Vampire
Vampire Bride
Nightstalker
Vampire
Vampire


1190/1190
4315/4315
Jewelled Bac...
Demon Back...
Moon Backpack
Golden Backp...
Minimap
Inventory
Combat Controls
General Controls
Logout Character
Change Character
Options
Quest Log
Character Profile
VIP List
Battle List
Dragon Back...
Vampire Bride
Vampire
Amaro De Quester
62
Using one of 370 great mana potions...
Local Chat
Server Log
English Chat
Advertising
15:12 Amaro de Quester [209]: exevo gran vis lux
15:12 Amaro de Quester [209]: exevo vis hur
15:12 Amaro de Quester [209]: exori flam
15:12 Amaro de Quester [209]: exori flam
15:12 Amaro de Quester [209]: exori max flam
15:13 Amaro de Quester [209]: exori frigo
15:13 Amaro de Quester [209]: exevo gran mas vis
15:13 Amaro de Quester [209]: utamo vita
15:13 Amaro de Quester [209]: exori vis
15:13 Amaro de Quester [209]: exori gran vis
15:13 Amaro de Quester [209]: exori flam
Say
Chat On
Set 6

Some Big Numbers

- ~150,000,000 page impressions per month
- ~20 terabyte web traffic per month
- ~55 terabyte game traffic per month
- 77 game worlds
- ~1,200,000 game logins per day
- ~500,000 different characters per day
- ~300,000 different accounts per day
- 95,000 active monthly subscriptions

Some Small Numbers

- People working on Tibia (in average)
  - 3 product managers
  - 4 programmers
  - 4 game content designers
  - 1 graphic artist (2D)
  - 2 software testers
  - 3 system administrators
  - 3 community managers
  - 9 customer support representatives

Architecture

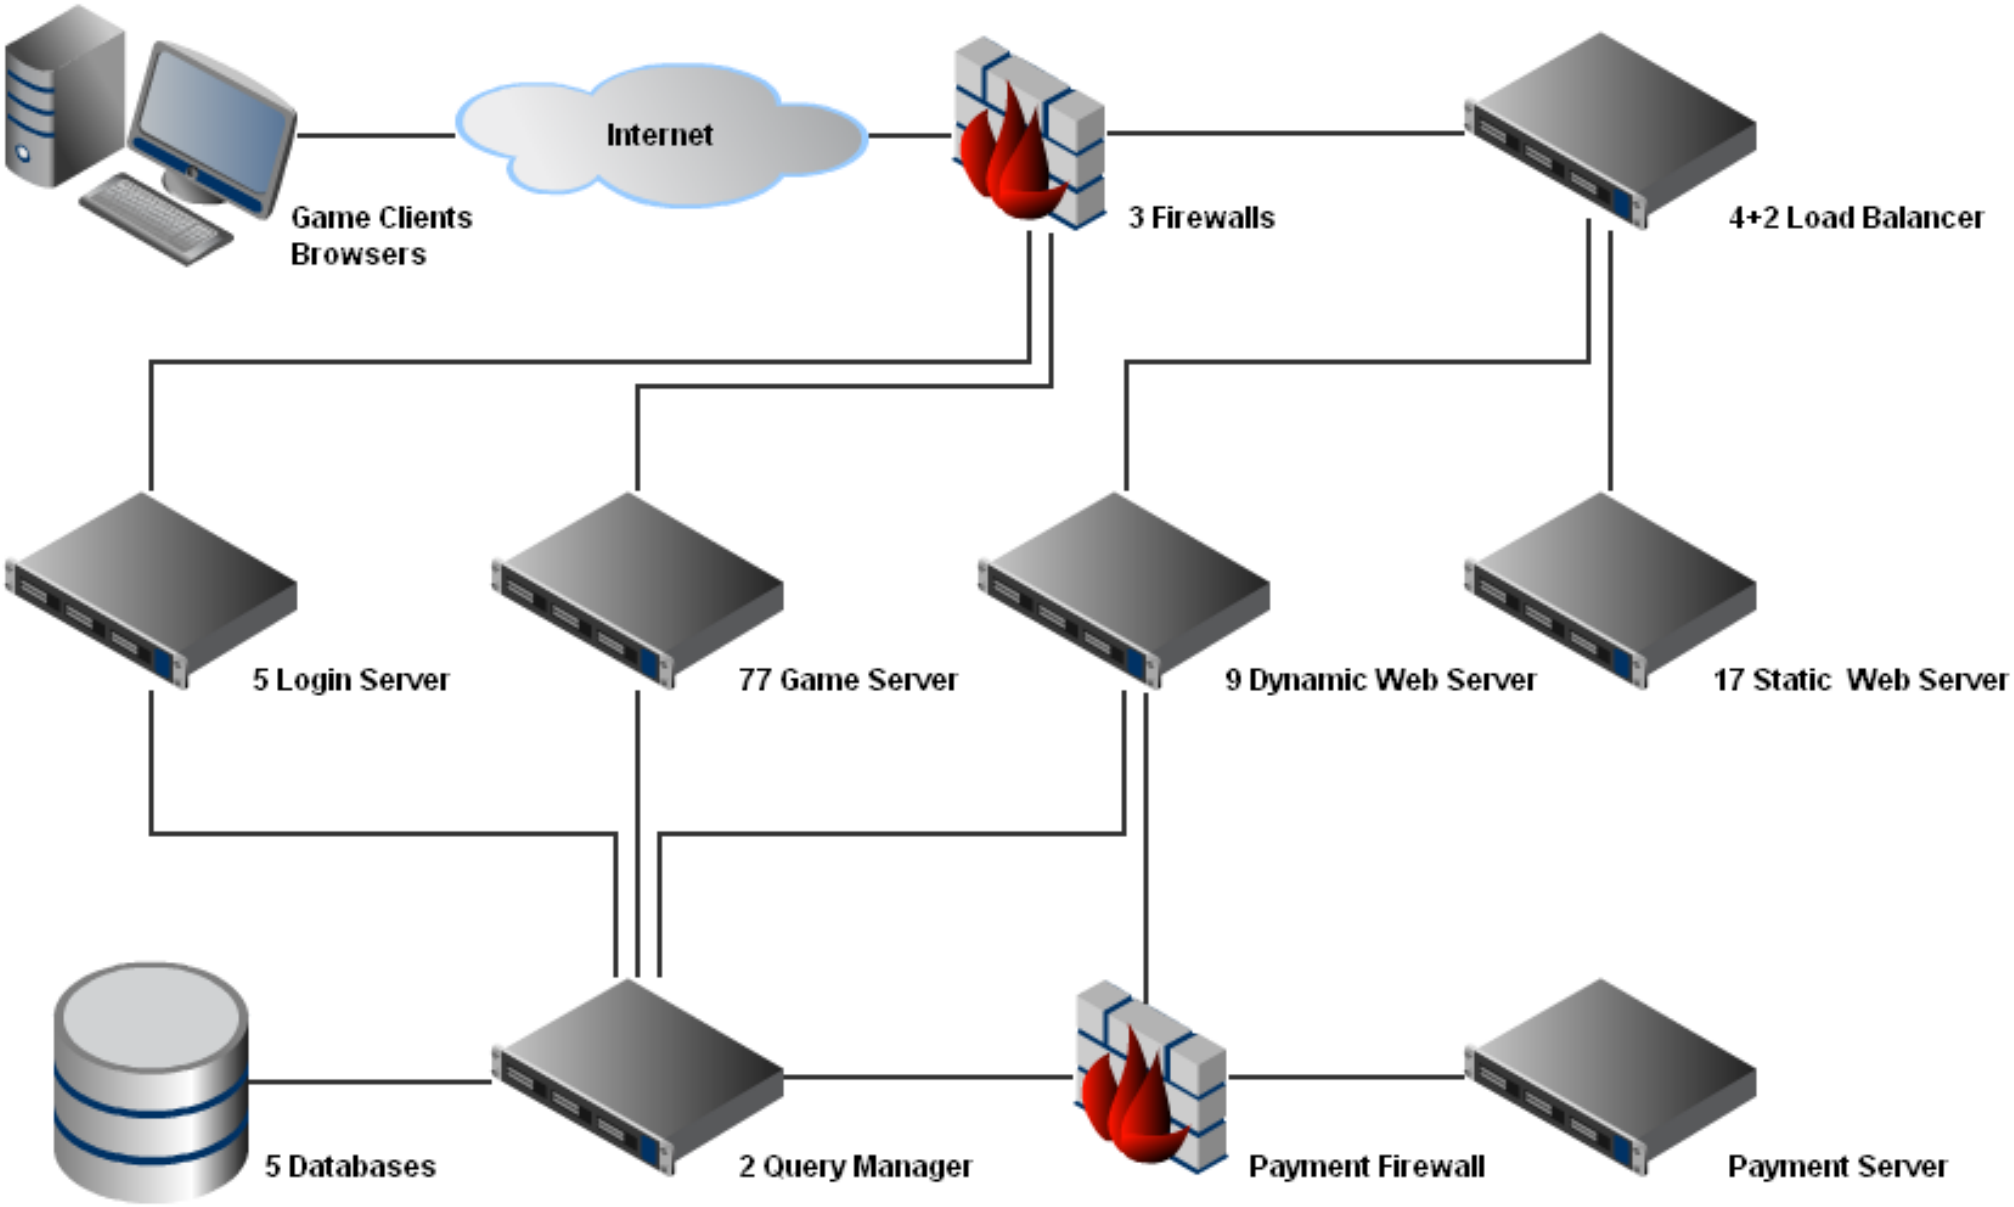

Servers: Default Hardware

- IBM BladeCenter
  - 2 power supplies
  - 2 network switches
  - 2 huge fans
  - 14 blade servers
    - 2 cores at 2.5 GHz
    - 4 GB ECC RAM
    - 2 hard disc 70 GB each in RAID
    - 2 network cards
  - CentOS 5.6

Servers: Locations

- own servers in Germany
  - 4 BladeCenter in Frankfurt
    - near DECIX
  - 1 BladeCenter in Nuremberg
    - near office in Regensburg
- rented servers in USA
  - hardware requirements similar to BladeCenter
  - in Houston and Dallas
    - near North and especially South America
- some spare blade servers as reserve
  - online but unused

Databases

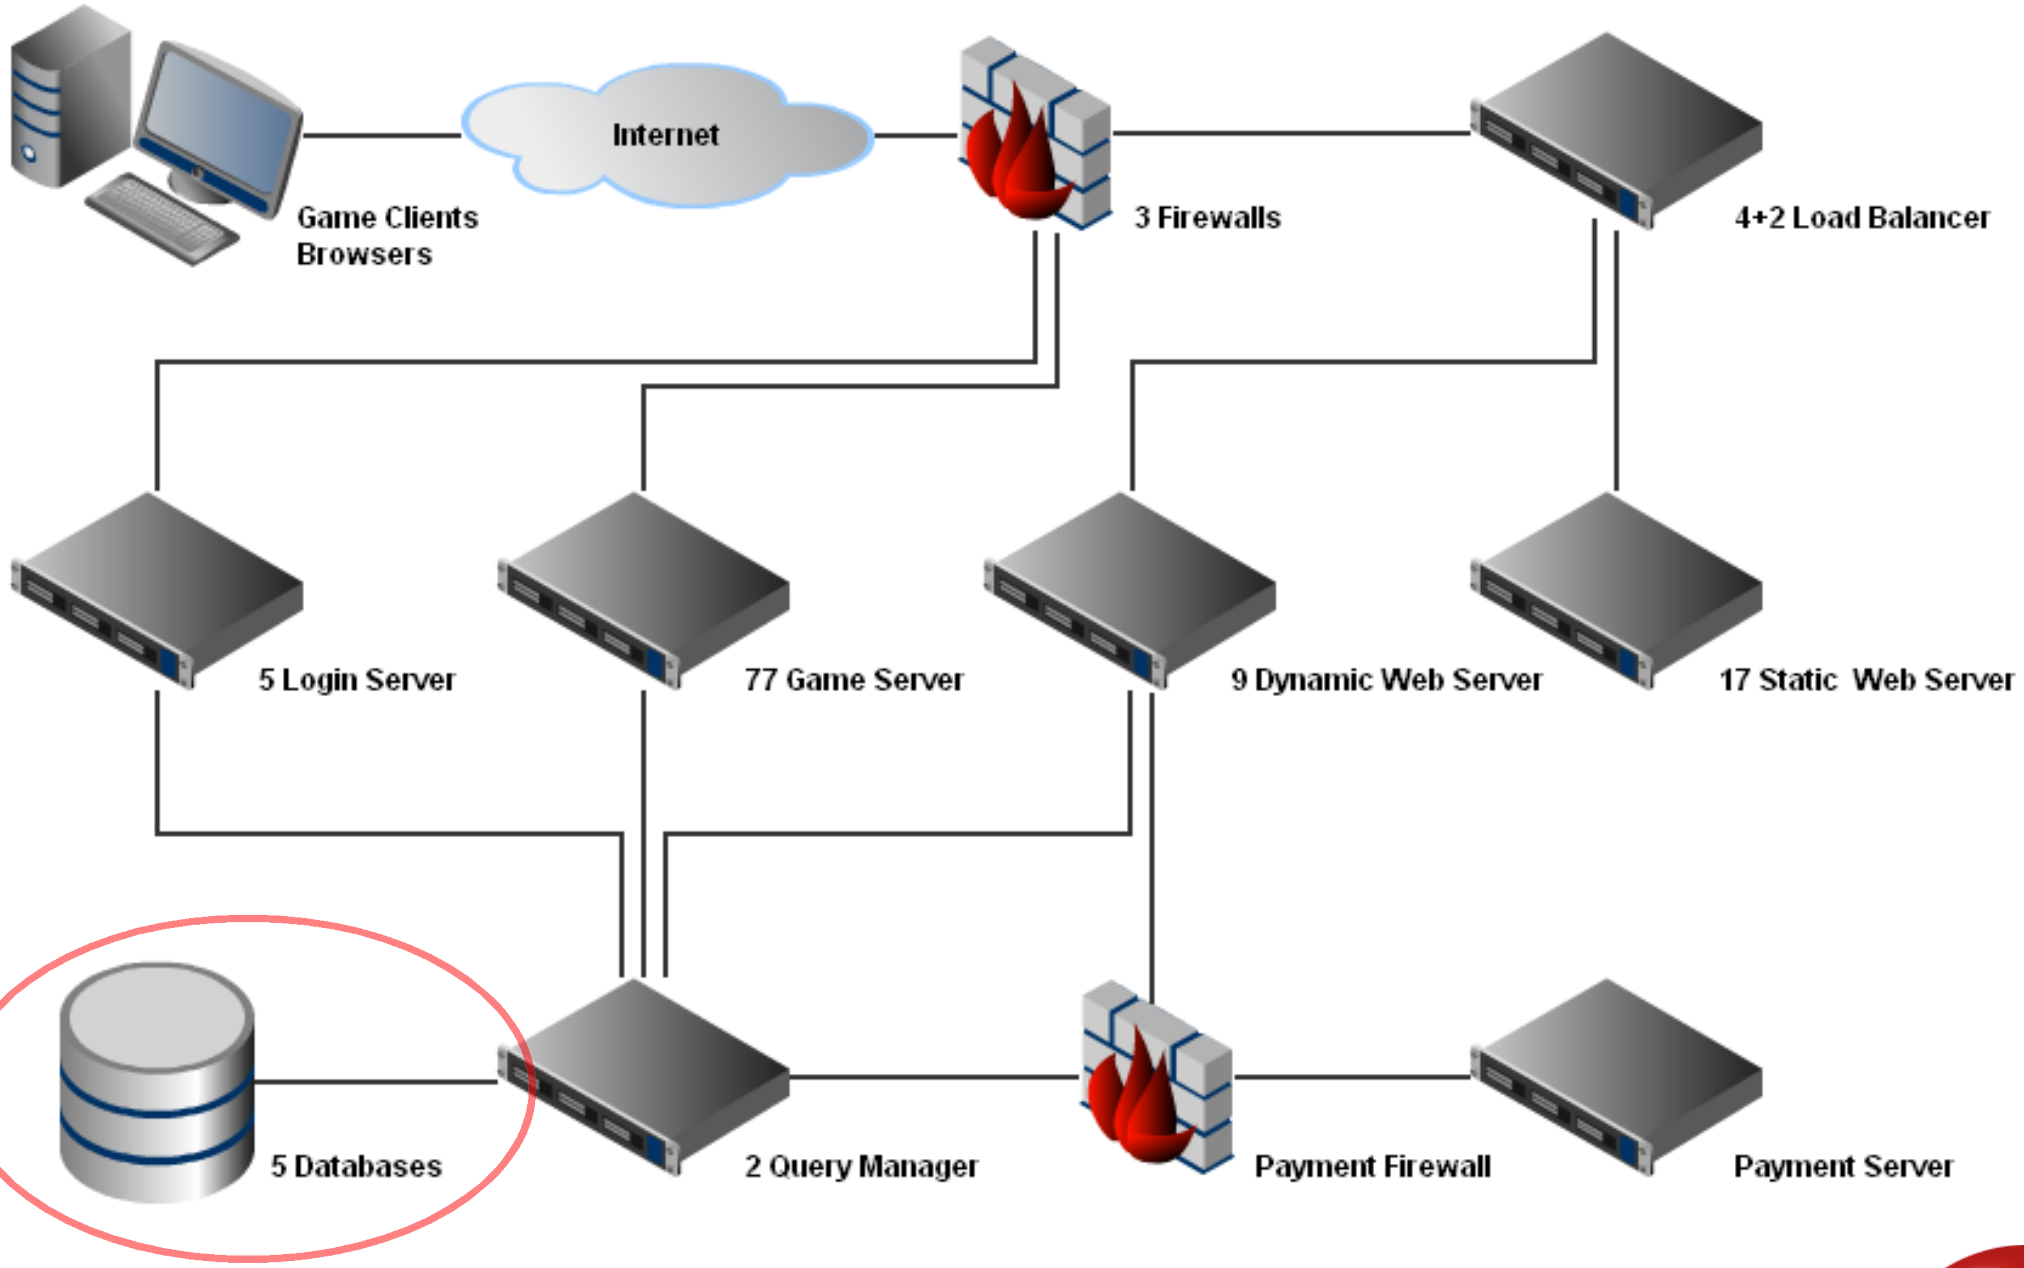
Internet
Game Clients
Browsers
3 Firewalls
4+2 Load Balancer
5 Login Server
77 Game Server
9 Dynamic Web Server
17 Static Web Server
5 Databases
2 Query Manager
Payment Firewall
Payment Server

Databases: Hardware and Software

CipSoft online entertainment

- one big database
  - 24 cores at 2.4 GHz
  - 128 GB ECC RAM mirrored (64 GB RAM)
- four smaller databases
  - 8 cores at 2.9 GHz
  - 24 GB ECC RAM
- all of them
  - storage area network
  - CentOS 5.6
  - PostgreSQL 8.4
  - no clustering, no mirroring
- located in Nuremberg

Databases: Data

- one big database
  - all account data
  - partial copy of character data
- four smaller databases
  - website data
    - statistics, etc.
  - volatile data
    - "who is online" list, etc.
  - management data
    - server lists, IP addresses, etc.
  - forum data

Databases: Software Choice

CipSoft online entertainment

- do not guess database performance, measure it!
- with realistic-as-possible data
  - structure
  - size
- we measured in 2005
  - copy of data and recorded requests from live system
  - PostgreSQL 7 vs Oracle RAC vs IBM DB2
  - PostgreSQL was slightly faster and a lot cheaper
  - reasons
    - all data in RAM (back then 6GB, now 25 GB)
    - 90% simple read operations (SELECT)

Query Managers

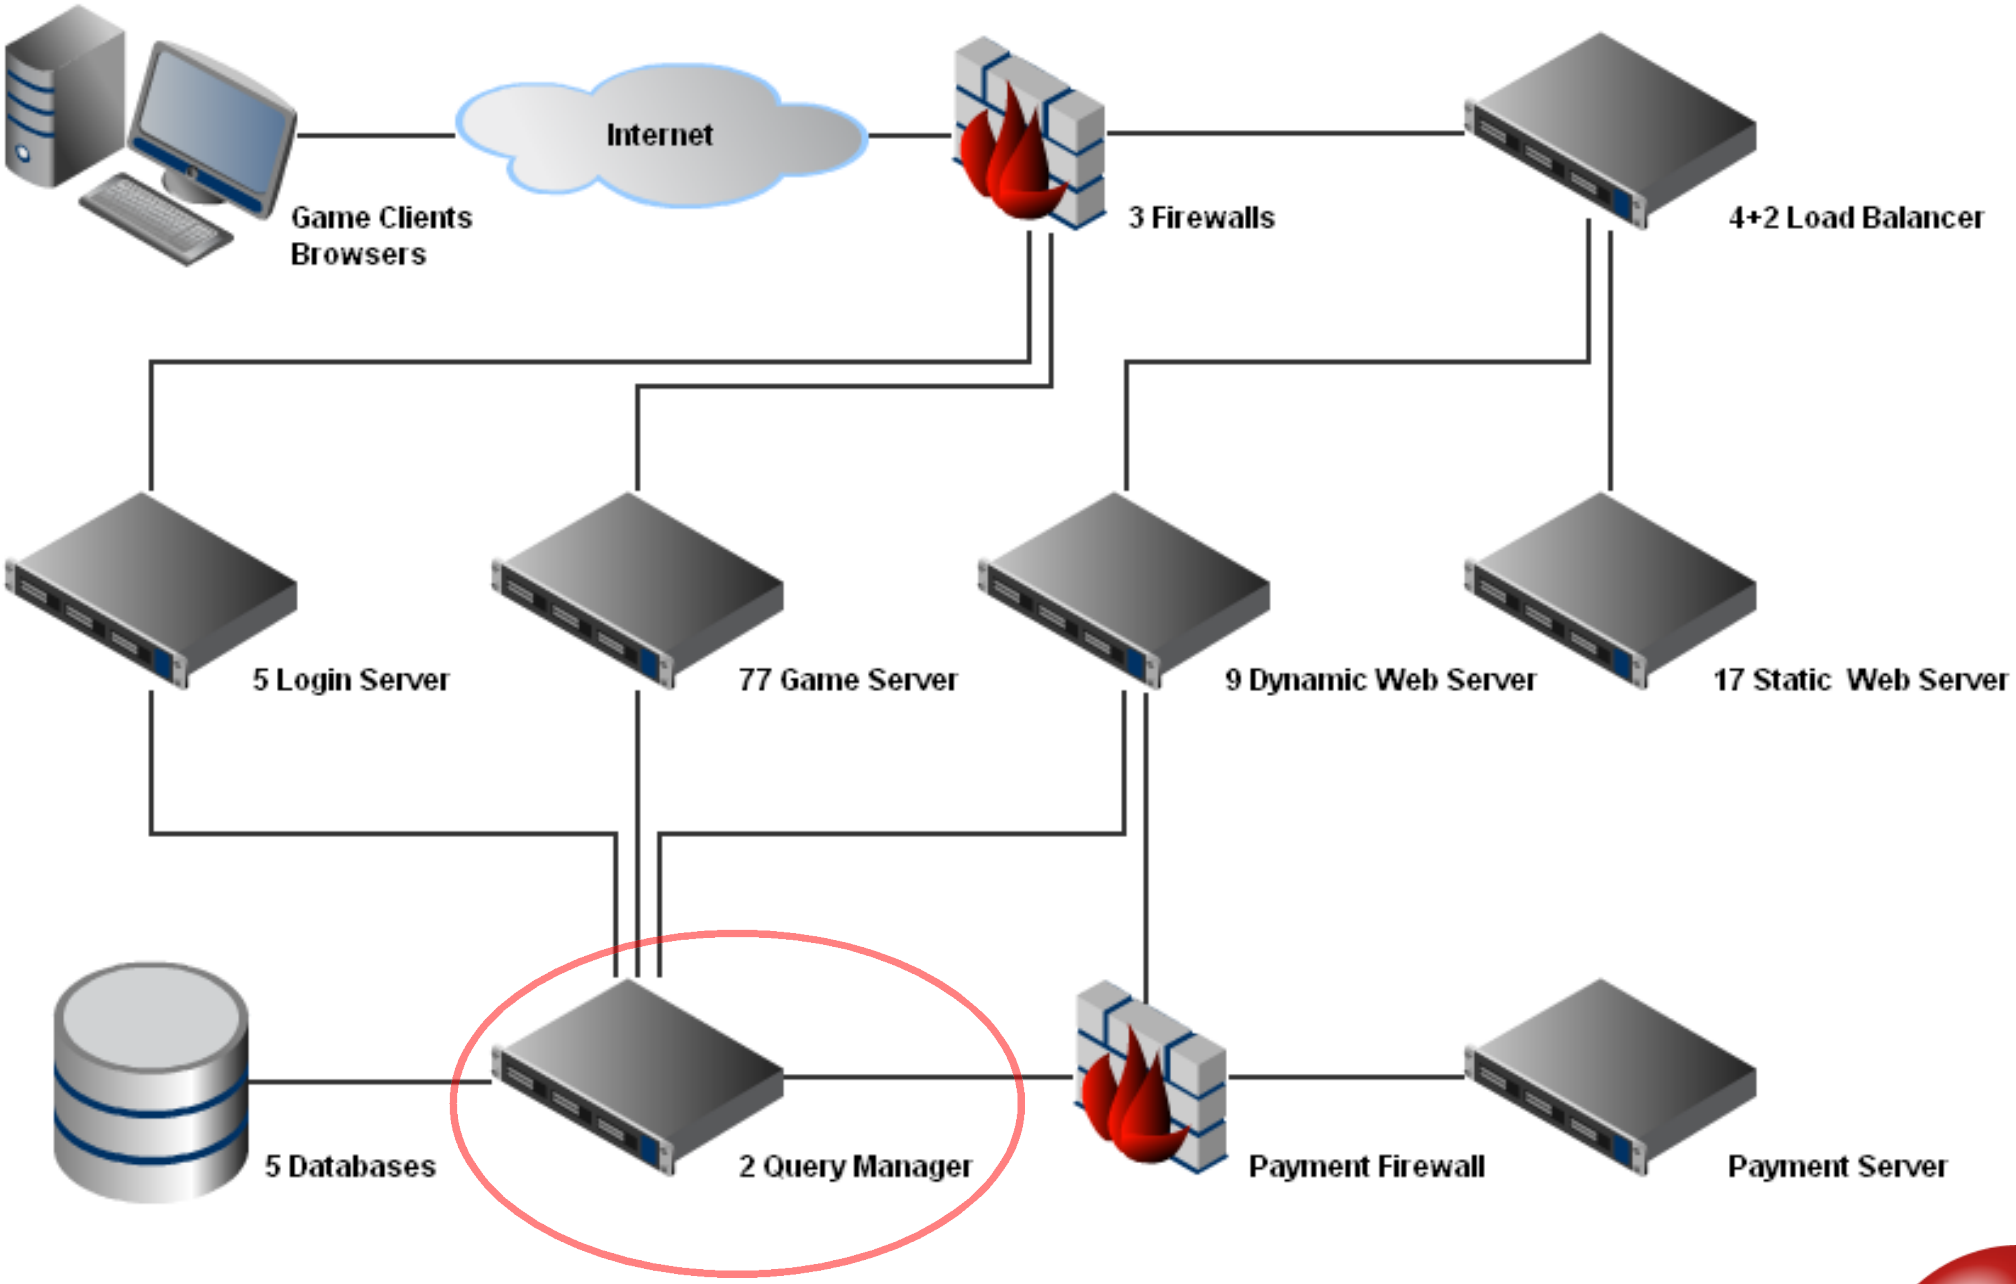

Query Managers

- custom server software
- intermediate layer in front of databases
- 2 of them
- physically right next to databases

Query Managers: Advantages

CipSoft online entertainment

- faster processing of requests from other servers
  - there is the Atlantic Ocean (150+ ms)
  - sometimes several SQL queries for request
  - sometimes C++ based logic for request
  - query managers physically right next to databases
- hiding data allocation
  - stores data in appropriate database
    - other servers don't care
  - simulates distributed database
  - not easily possible with PostgreSQL

Query Managers: Advantages

- additional access control
  - no direct access from web servers to database
  - no commodity software
  - defined requests with strict syntax
  - different access rights for different servers
    - web server
    - game server
    - payment server
- profiling
  - count types of requests
  - measure times of requests

Query Managers: Disadvantages

CipSoft online entertainment

- yet another layer
  - implementation
  - testing
  - administration
  - point of failure
- limits
  - amount of connections
  - amount of requests
  - etc.

Query Managers: Connections

- opening connections to all databases at startup
- accepting connections from other servers
  - TCP/IP
  - SSL encrypted
  - proprietary binary protocol

Query Managers: Code

- written in C++
  - 30,000 LOC (lines of code)
  - 5,500 LOC Tibia's shared code
  - 28,000 LOC CipSoft's network and utility library
- SQL statements only in this server
- prepared queries wherever possible
- stateless (after authorization)
- multithreaded

Game Servers

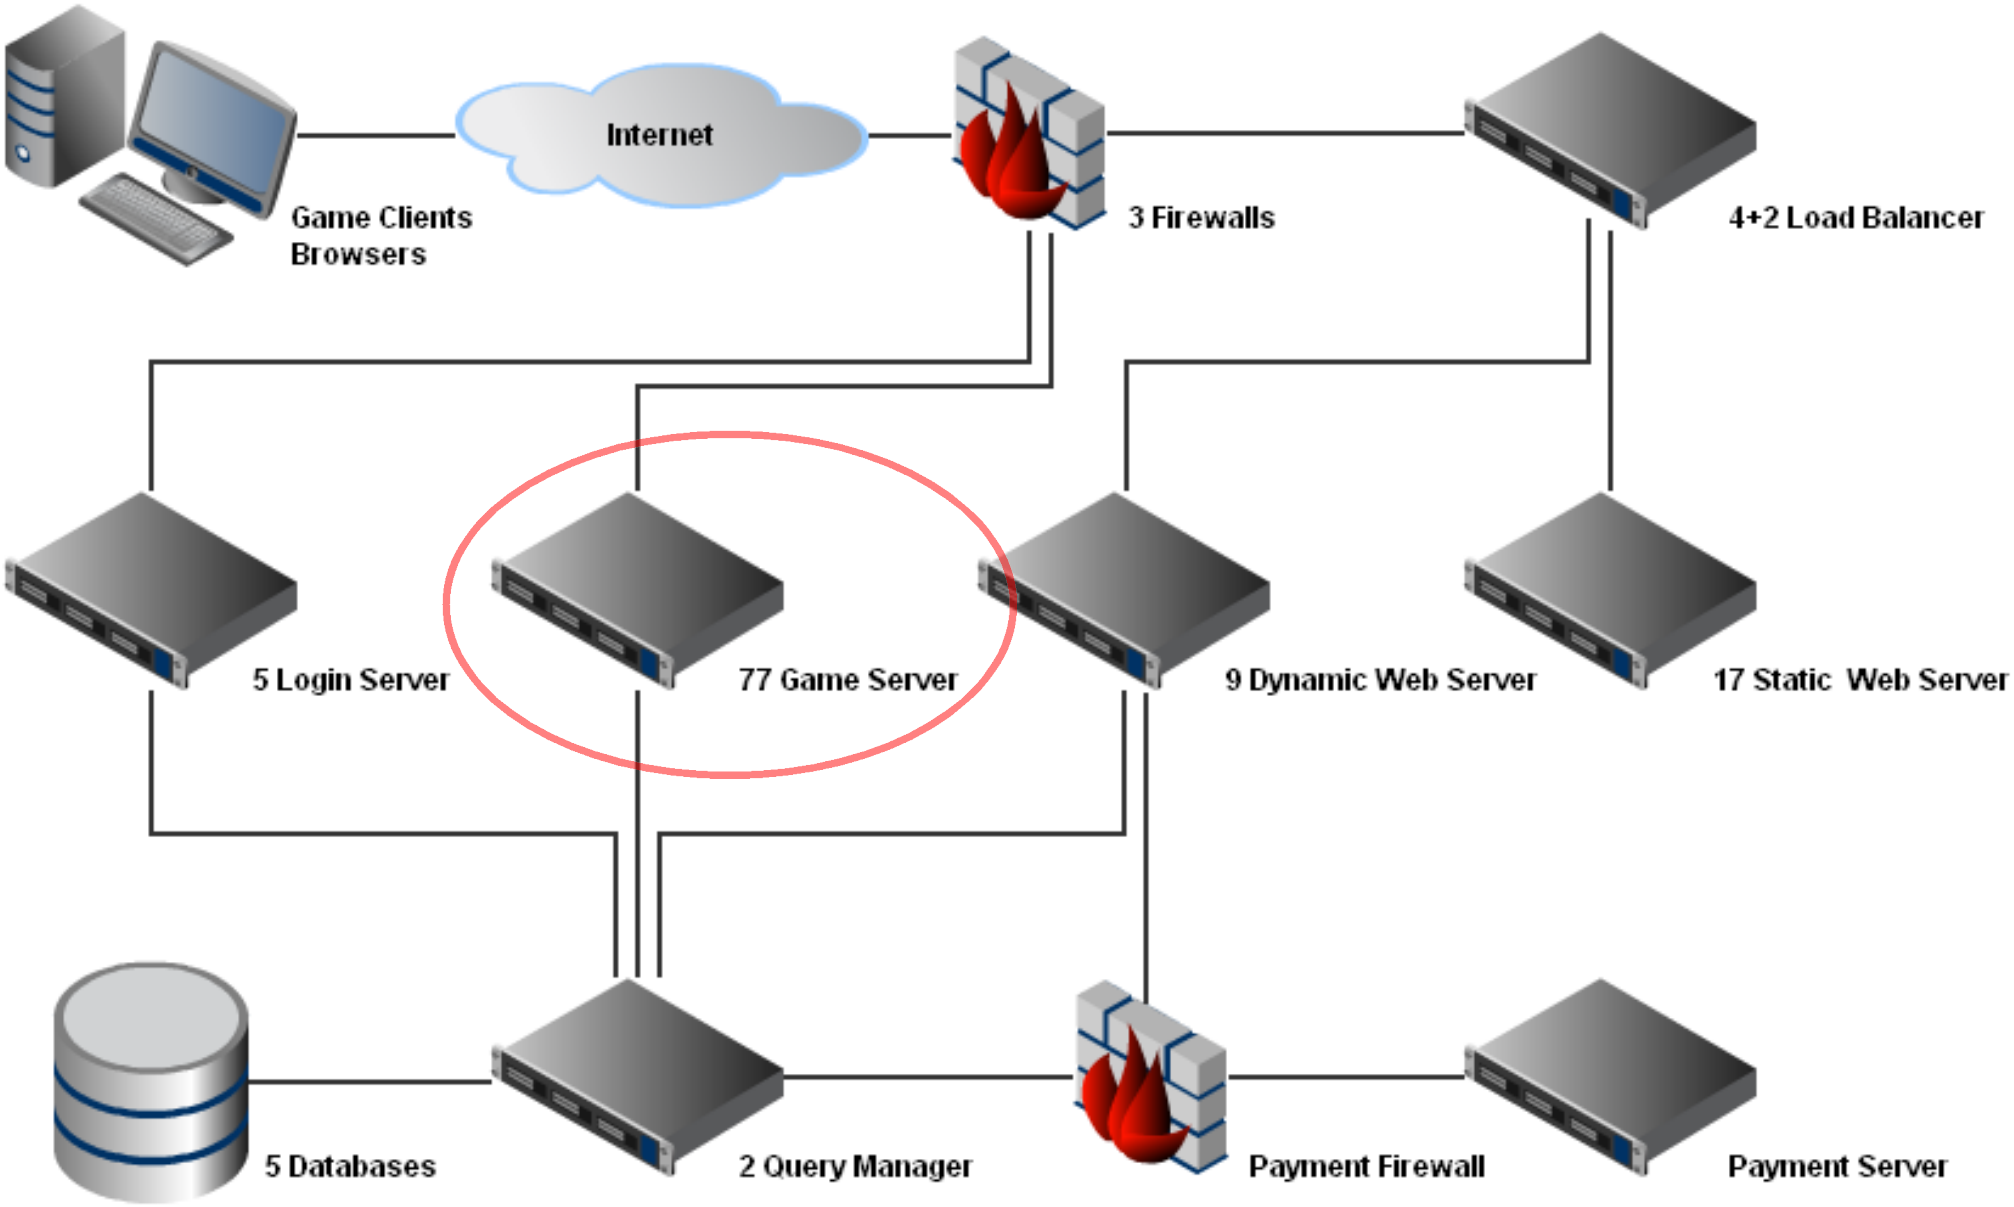

Game Servers

- 1 game world runs on 1 blade server
- 77 game worlds
  - half located in Frankfurt
    - near to DECIX
  - half located in Dallas
    - near to North and South America
- simulation of the game world
- maximum of 1050 characters online
  - formerly restricted by CPU load
  - currently restricted by game world size
    - game design decision

Game Servers: Data Distribution

- account data in database
- character data local on hard disc
  - one (proprietary) text file per character
  - some of it copied into database for use on website
  - loaded on demand (character login)
  - daily backup
- world data local on hard disc
  - ~1,700 (proprietary) text files for definitions (~15 MB)
  - ~17,500 (proprietary) text files for world map (~300 MB)
  - same again for "current" version of world map
  - everything loaded at game server startup
  - daily backup

Game Servers: Connections

- opening 10 connections to query managers at startup
  - TCP/IP
  - SSL encrypted
  - proprietary binary protocol
- accepting connections from clients
  - TCP/IP
  - RSA encrypted login request
  - XTEA encrypted afterwards
  - proprietary binary protocol

Game Servers: Code

- written in C++
  - 45,000 LOC
  - 5,500 LOC Tibia's shared code
  - 28,000 LOC CipSoft's network and utility library
- multithreaded...
- ...except the whole world simulation

Game Servers: Code

- origin of world simulation in age of single CPU core
- advantage
  - no synchronization within world simulation
- disadvantages
  - does not scale
  - limited by performance of one CPU core
- the plan so far
  - keep world simulation as it is
  - offload anything else in supporting threads
  - think about it for the next game...

Game Servers: Code

- supporting threads
  - acceptor/receiver/sender threads
    - epoll, edge triggered, BSD sockets
    - efficient on Linux
    - not efficient when using OpenSSL
    - default model in our network library
    - our solution, there are others
      - Google "The C10K Problem"
  - reader/writer threads
    - main thread shall not block on hard disc i/o
  - RSA decryption thread
    - intentional bottleneck against denial of service attacks on CPU

Login Servers

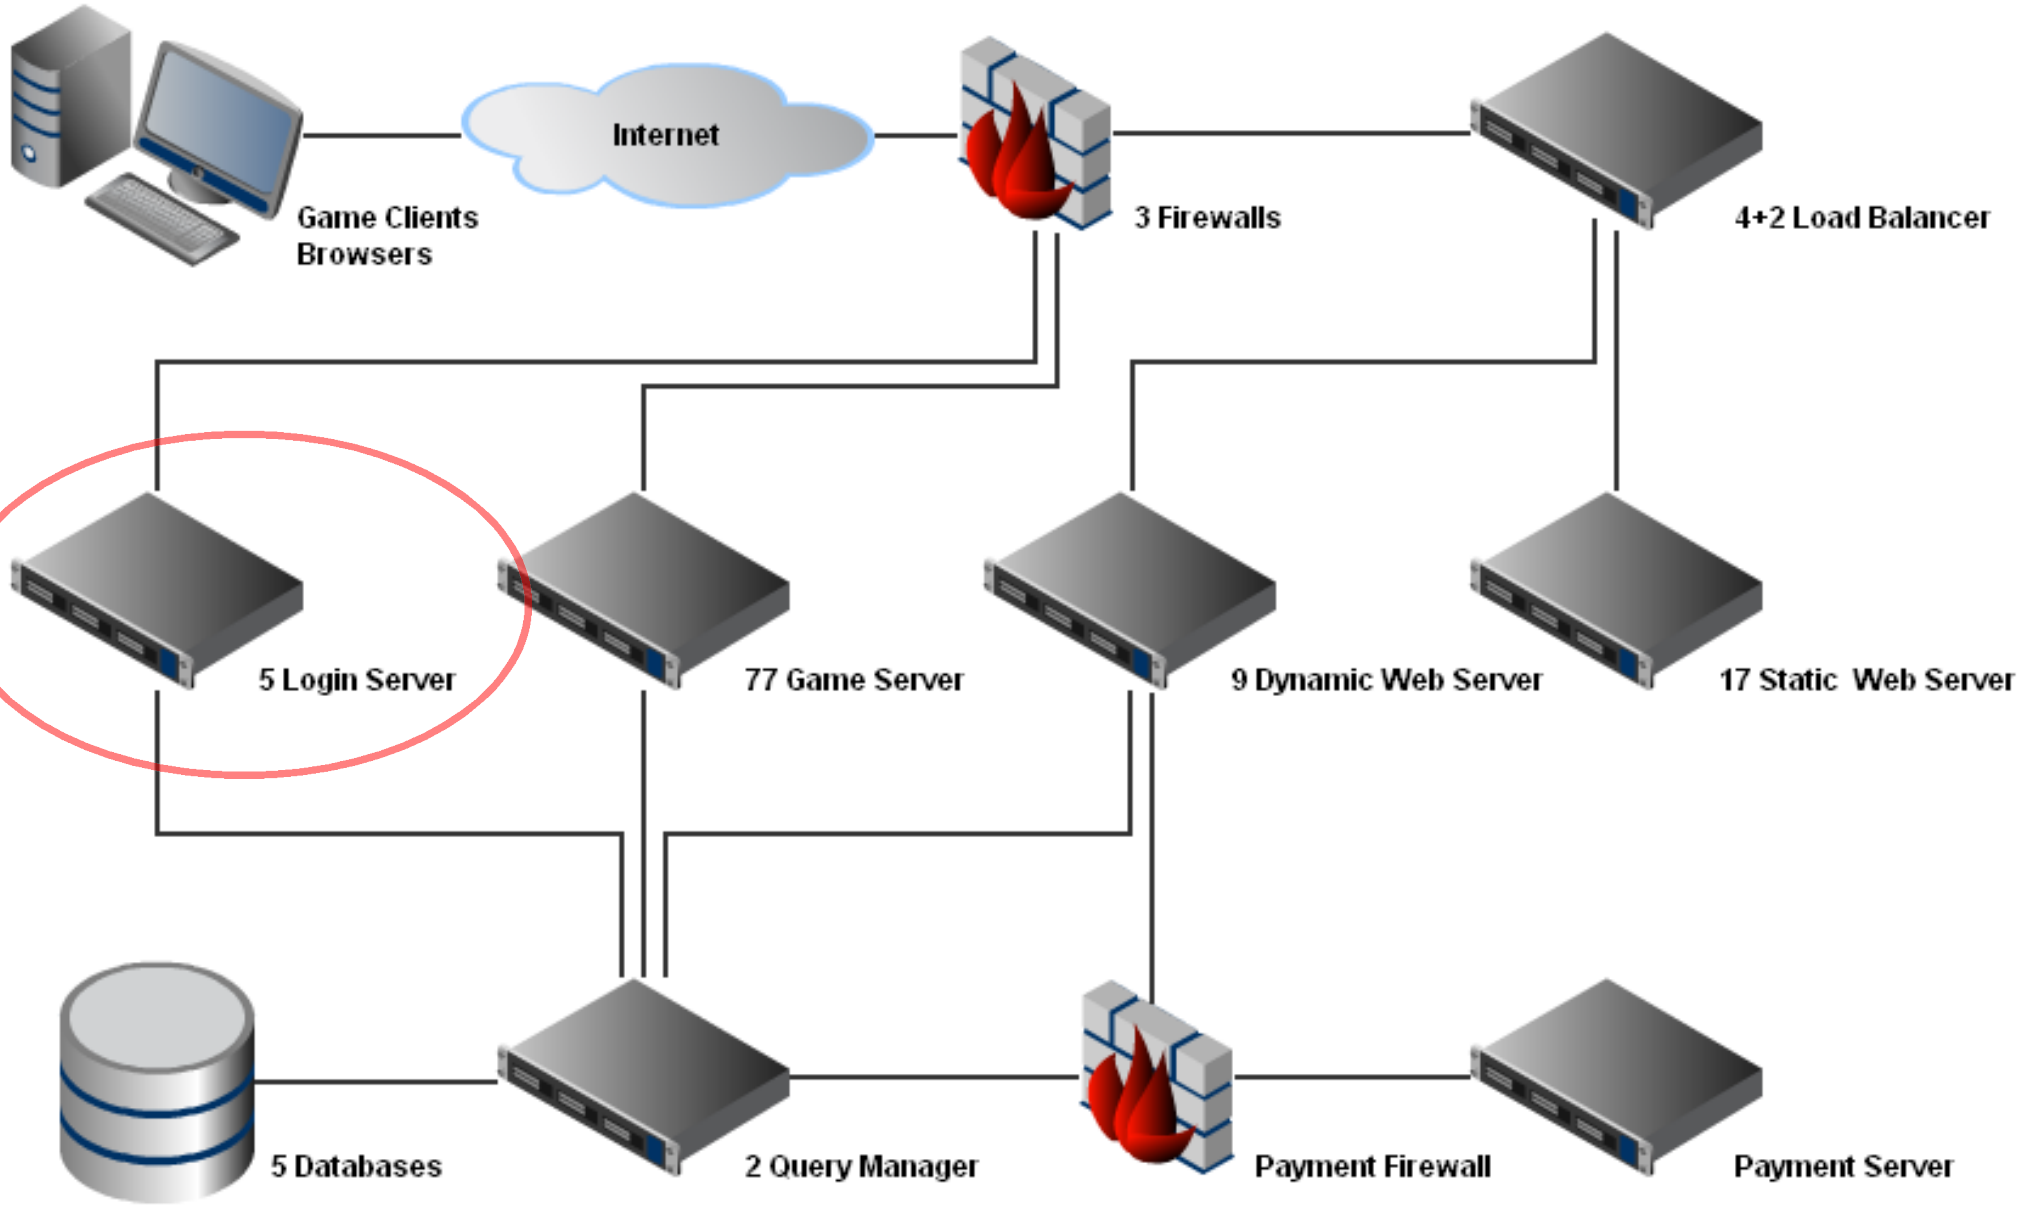

Login Servers

- custom server required for stand-alone client
  - client update
  - account authentication
  - character selection
  - guidepost towards game servers regarding IP addresses
- 5 of them
  - 1 in Nuremberg
  - 2 in Frankfurt
  - 2 in Houston
- 10 DNS entries
  - in 2 domains (login01.tibia.com, tibia01.cipsoft.com, etc.)
  - hardcoded in clients

Login Servers: Connections

- opening 10 connections to query managers at startup
  - TCP/IP
  - SSL encrypted
  - proprietary binary protocol
- accepting connections from game clients
  - TCP/IP
  - RSA encrypted login request
  - XTEA encrypted afterwards
  - (simple) proprietary binary protocol

Login Servers: Code

- written in C++
  - 6,000 LOC
  - 5,500 LOC Tibia's shared code
  - 28,000 LOC CipSoft's network and utility library
- stateless
- multithreaded

Game Clients

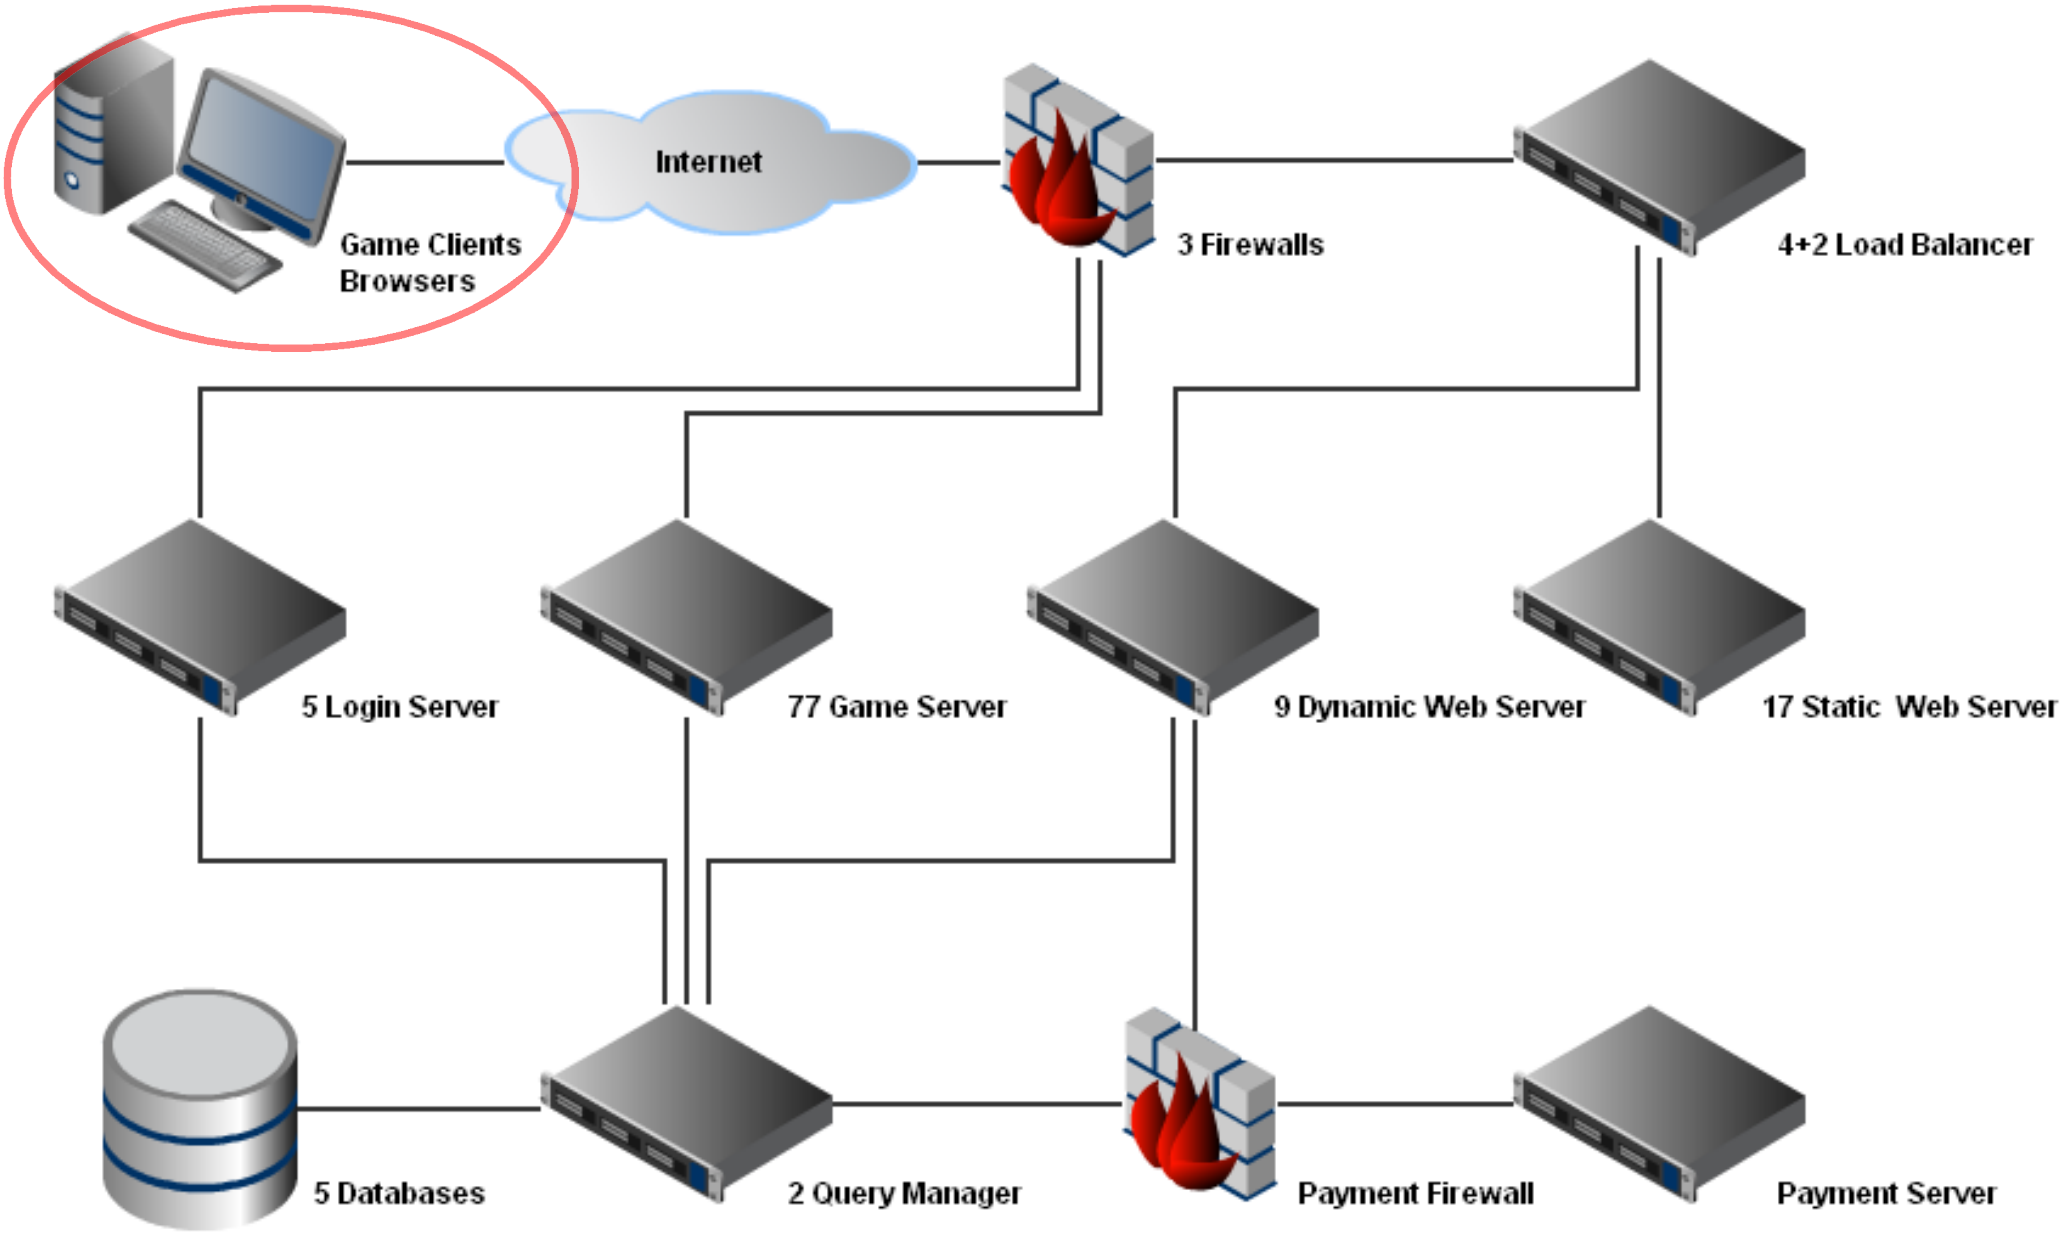

Game Clients: Stand-Alone Client

- Windows XP / Vista / 7
- Windows 95 / 98 / ME / 2000 until July 2011
- Linux
- 27 MB installer
- automatic update (over login server)
- storing data on hard disc
  - object definitions and images: 50 MB
  - discovered mini map: up to 200 MB
- written in C++
- 63,600 LOC
- single threaded

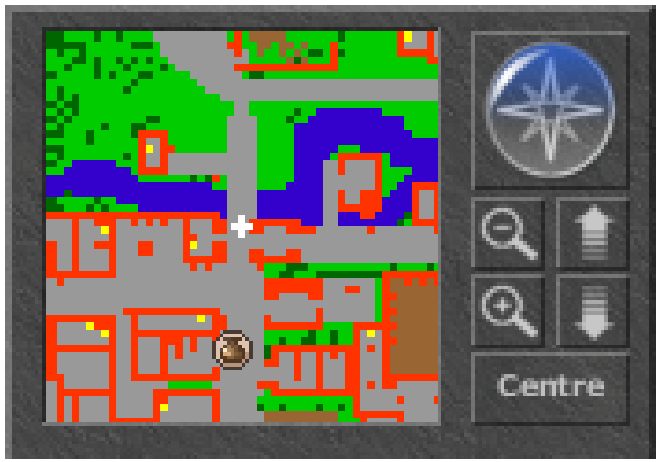

Game Clients: Flash Client

- browser based client
- 1.5 years of development
- available since June 2011
- still has "Beta" label
- automatic update (over web servers)
- caching data in Flash cookies
  - object definitions and images: 40 MB
  - discovered mini map: up to 200 MB
- written in ActionScript3
- 66,000 LOC and growing
- single threaded

Game Clients: Connections

CipSoft online entertainment

- opening 1 connection...

  ...first to login server

  ...and later to game server
  - TCP/IP
  - RSA encrypted login request
  - XTEA encrypted afterwards
  - proprietary binary protocol

Encryption: RSA

- asymmetric encryption with RSA
  - well known algorithm
  - secure enough
  - open source implementation without dependencies
    - not OpenSSL library (too big)
  - 1024 bit key
  - public key hardcoded in game client
  - private key hardcoded in game server
  - used for login request
    - to login server
    - to game server

Encryption: XTEA

- symmetric encryption with XTEA
  - well known algorithm
  - secure enough
  - fast
  - open source implementation without dependencies
  - symmetric key
    - created by client
    - wrapped into login request
  - used for everything except login request

Encryption: ISAAC

- random number generation with ISAAC
  - secure enough
  - open source implementation without dependencies
  - never ever use `rand()` function for anything remotely related to encryption!

Encryption: Connection Handling

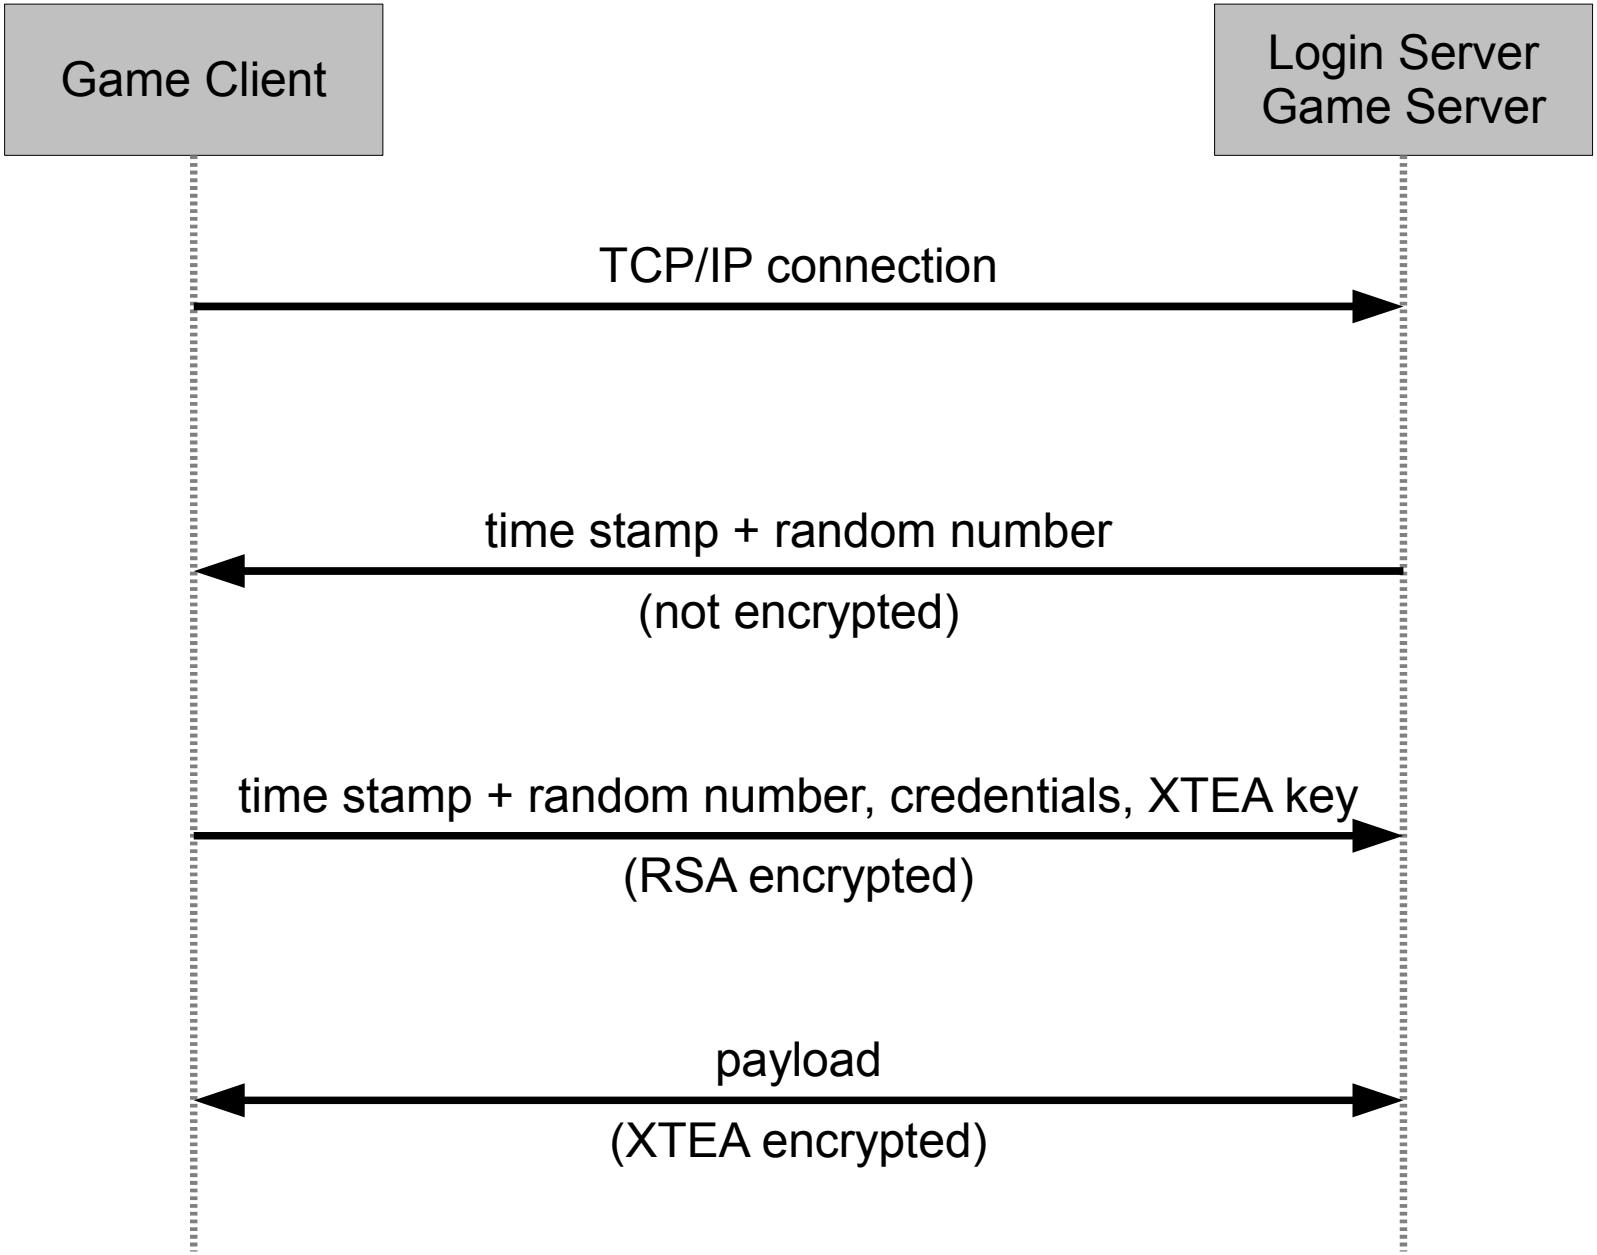

Encryption: How We Failed

- fail #1
  - used `rand()` function
  - got XTEA keys brute-forced
- fail #2
  - used no time stamp + random number
  - got attacks by replaying (MITM) recorded login packets
- fail #3
  - swapped `p` and `q` in server side implementation of RSA
  - got private key cracked
- conclusion
  - anything in encryption not 100% correct...
  - ...your whole encryption breaks

Web Servers and Load Balancers

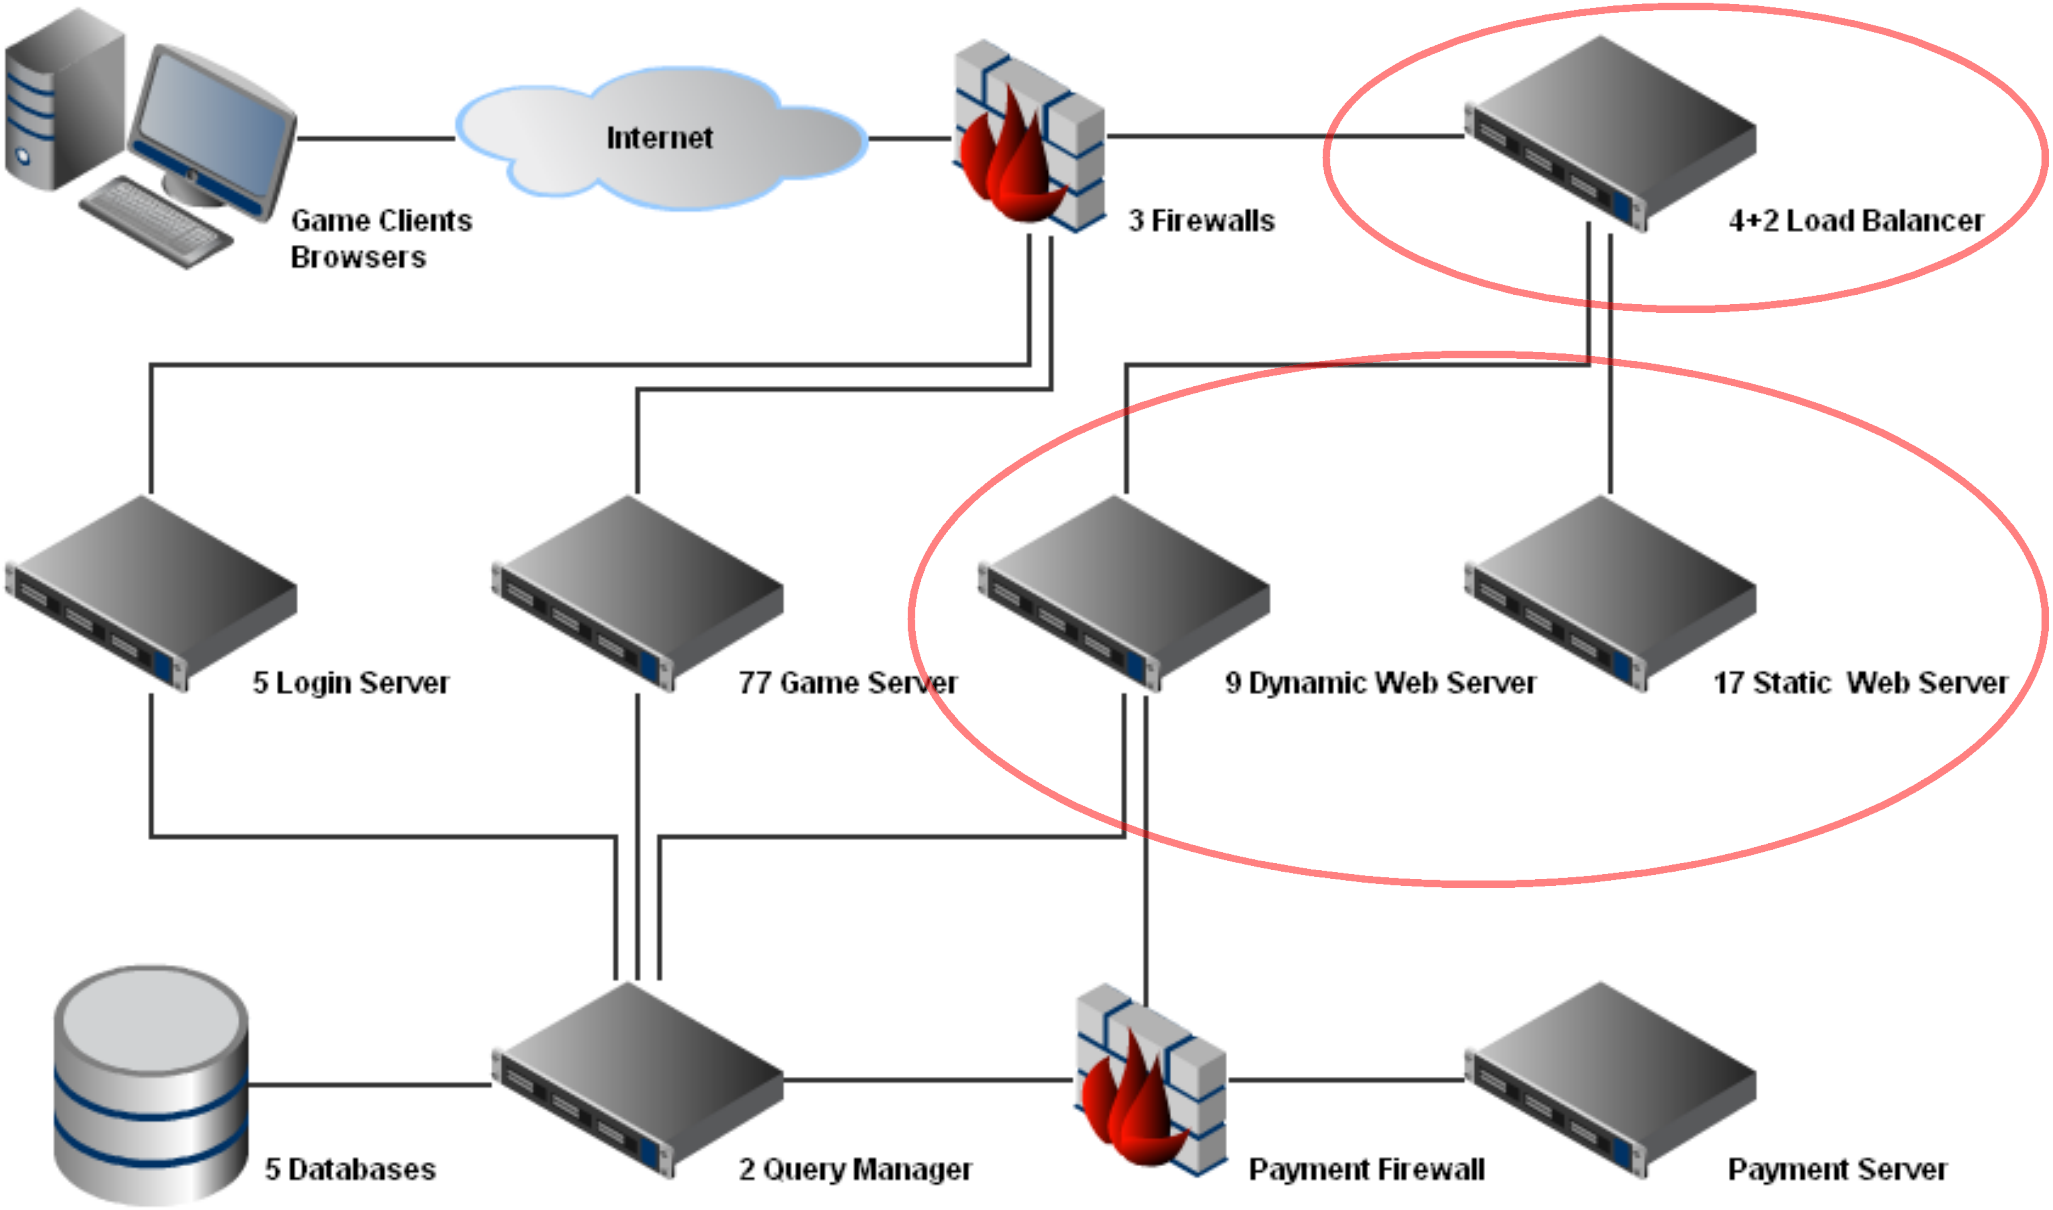

Web Servers and Load Balancers

- website
  - information
  - client downloads (stand-alone client)
  - client data (Flash client)
  - statistics
  - account management
    - account creation
    - character creation
    - guild management
    - house management
    - payment
  - forum

Web Servers and Load Balancers

- 17 static web servers
  - 13 http, 4 https
  - located in USA (cheaper web traffic)
  - Apache 2.2
- 9 dynamic web servers
  - 7 http, 2 https
  - located in Germany (near to databases)
  - Apache 2.2
  - PHP 5.3
- 6 load balancers
  - Linux Virtual Server

Content Delivery Network

- big change in April 2011
- no more static web servers
- only 2 load balancers left
- now using a content delivery network
  - hosts and caches all static web content
  - reroutes and caches all dynamic web content
  - Akamai

Content Delivery Network

- advantages
  - shorter load times of static web content for customers
  - no need for extra server capacity during peak times
  - better protection against DDoS attacks
  - all in all ~60% cheaper
    - less server rental costs
    - less administration costs

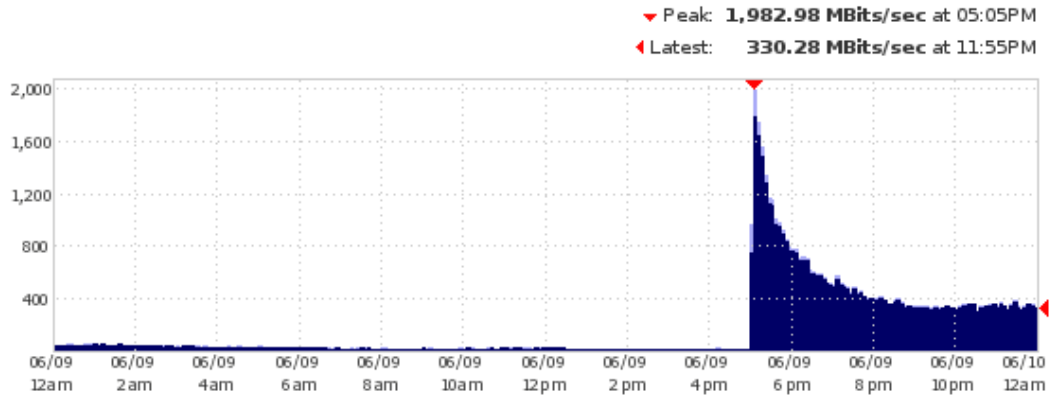

- disadvantages
  - initial setup (not that big)
  - their system, their rules
  - update of cached data not instant (obviously)

Payment Server

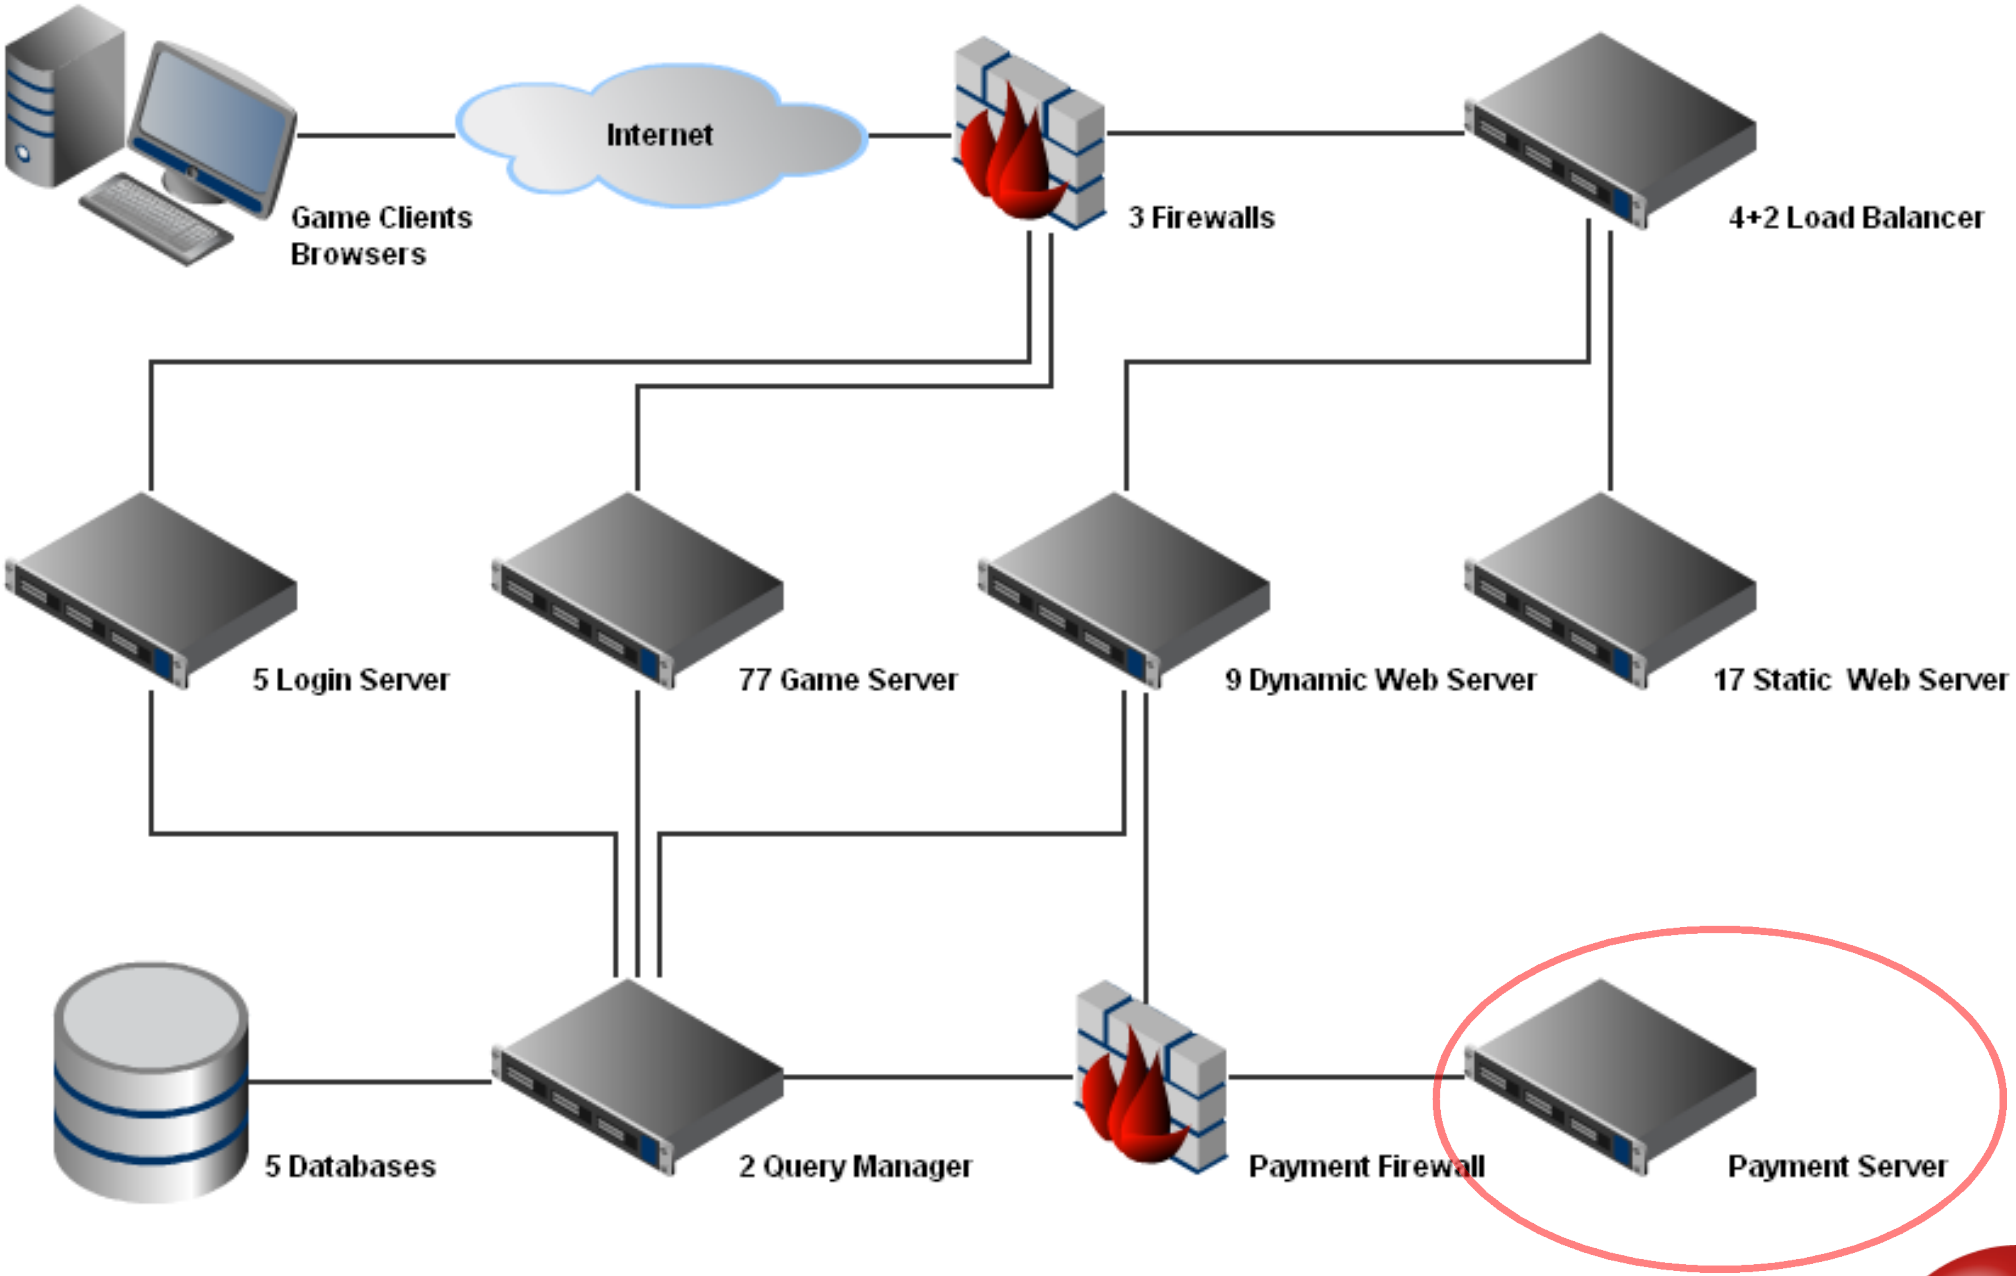

Payment Server

- handling data exchange with payment provider
- accepting connections from query managers
  - TCP/IP
  - SSL encrypted
  - proprietary binary protocol
- written in C++
  - 11,000 LOC
  - 5,500 LOC Tibia's shared code
  - 28,000 LOC CipSoft's network and utility library
- stateless (after authorization)
- multithreaded

Firewalls

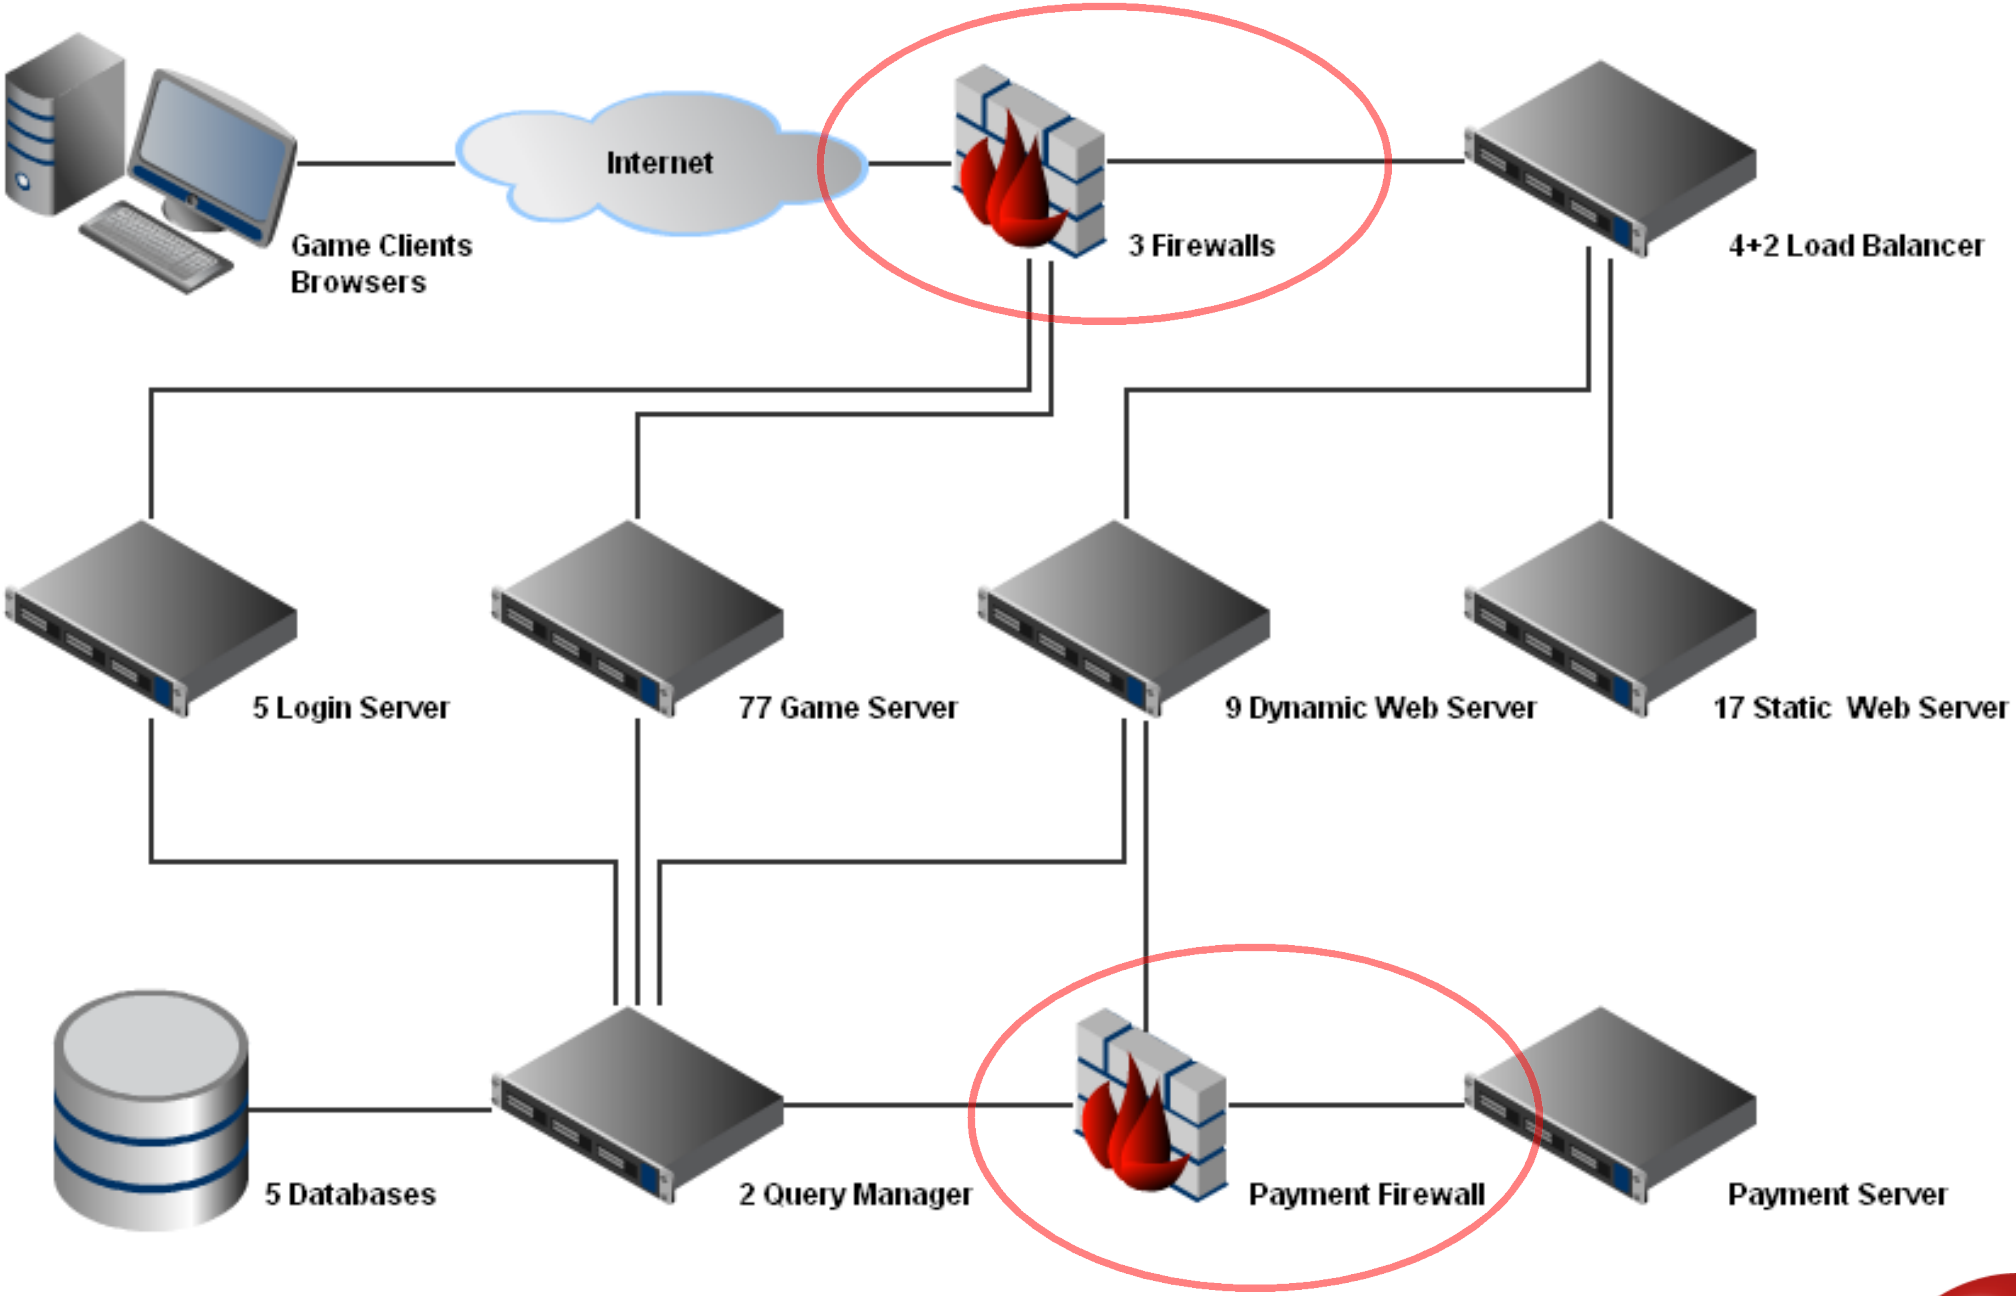
Internet
Game Clients
Browsers
3 Firewalls
4+2 Load Balancer
5 Login Server
77 Game Server
9 Dynamic Web Server
17 Static Web Server
5 Databases
2 Query Manager
Payment Firewall
Payment Server

Firewalls

- 3 big hardware firewalls
  - one for each location
    - Nuremberg
    - Frankfurt
    - Dallas/Houston
  - every server behind one of those
  - purpose: defence against packet rate DDoS attacks
- 1 small hardware firewall
  - in front of payment server
  - required for PCI-DSS
  - purpose: defence against web vulnerability attacks

Distributed Denial of Service Attacks

CipSoft online entertainment

- information known to users
  - list of users online (from website)
  - IP address of game server (after login)
- impact of DDoS attack on game server
  - interrupts connections
    - of all users of game server
    - of all users of datacenter (if big enough)
- but whyyyyy?
  - disconnect in Tibia = usually character death = XP loss
  - ingame conflicts
  - because they can

Lessons Learnt

- intended architecture improvements for next game
  - **better** resistance against DDoS attacks
    - by design, not just by bigger firewalls
  - better multithreading
    - no big, undivideable thread
  - better scalability
    - cloud style
  - well known formats instead of proprietary ones
    - XML
    - JSON
    - etc.

Next Game Architecture

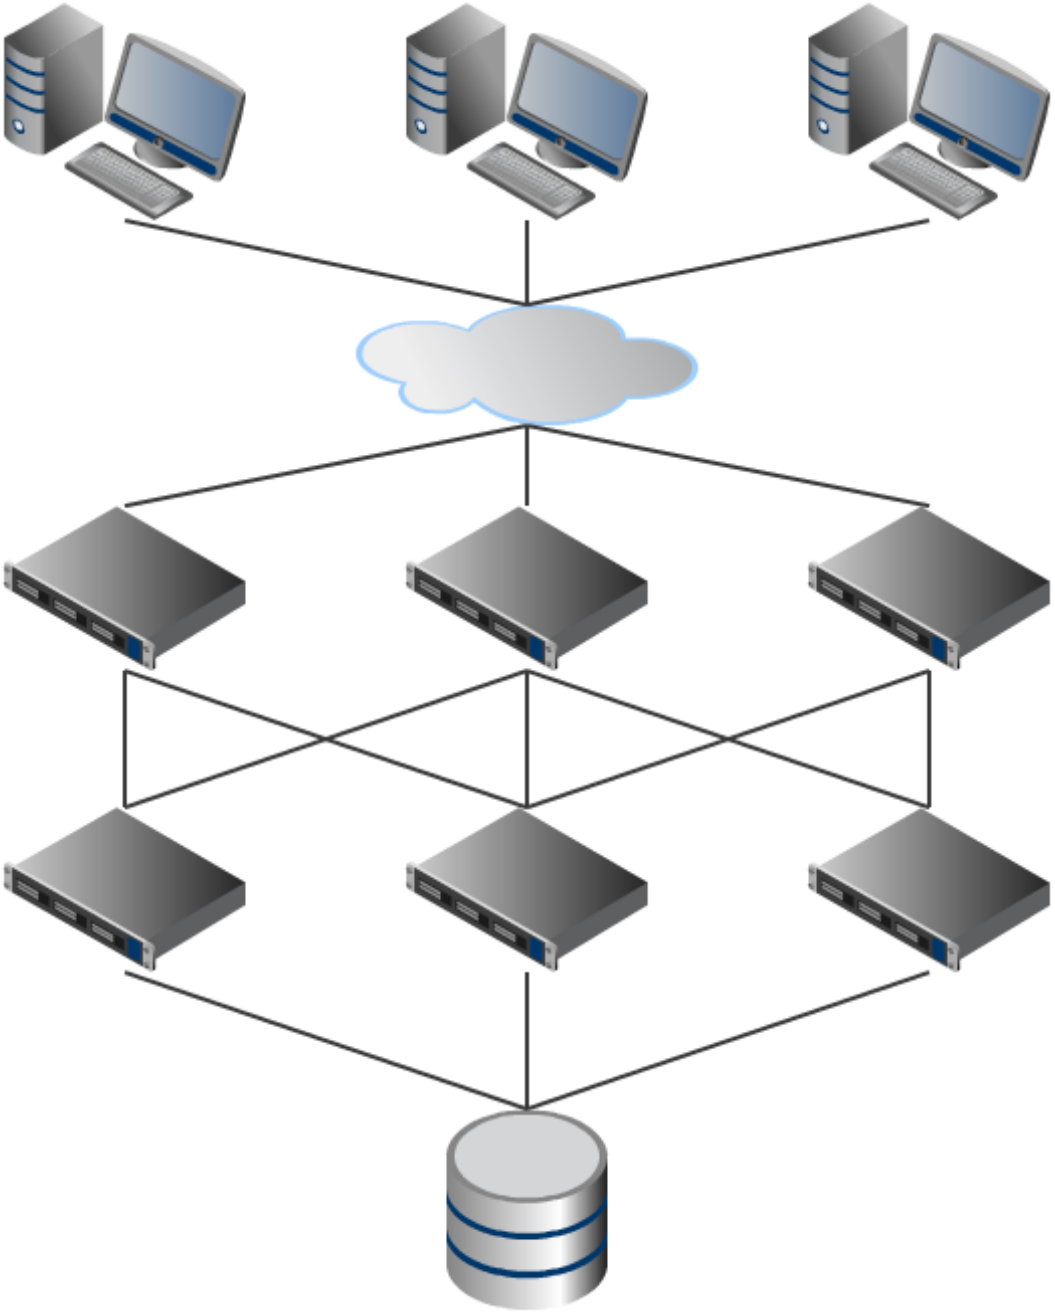

Game Clients

Internet

Dispatcher
reachable from Internet

Game Servers
not reachable from Internet

Database
not reachable from Internet

Next Game Architecture: Advantages

CipSoft online entertainment

- DDoS attack on dispatcher less harmful
  - no direct impact on game servers
  - disconnects unknown group of users
    - "unknown" is the big advantage
    - the more dispatchers the less impact
  - disconnected users have only small drawback ingame
    - game design related
  - disconnected users can instantly reconnect using any other dispatcher
- dispatchers and game servers could be in the cloud
- dispatchers could be run on plenty locations worldwide

Next Game Architecture: Disadvantages

CipSoft online entertainment

- more layers
  - more implementation, testing, administration
  - more latency
- independency of game servers required for scalability
- game design restrictions
  - latency must not be that important
  - disconnect must not be that painful

Thanks!

- rudy@cipsoft.com
- http://www.cipsoft.com
- http://www.tibia.com

We're hiring!
gamescom
"Jobs & Karriere"
hall 8 booth B40
http://www.cipsoft.com/jobs