

Life Sciences at Genentech

(and IT infrastructure)

(and Databases)

John “Scooter” Morris, Ph.D.

Genentech, Inc.

Introductions

Who are you?

- IT Professional?
- Scientist?
- Business?
- Wrong place?

Outline

Part 1

- The Company

Part 2

- IT Infrastructure

Part 3

- Architecture Project
- Unified Database Tier

The Company

Genentech – Mission

“Genentech is a leading biotechnology company that discovers, develops, manufactures and commercializes biotherapeutics for significant unmet medical needs.”

Significant unmet medical needs

- At the end of the day – it’s about helping people...



Genentech – 5x5

Genentech's corporate strategy

Become the world's (leading) biotechnology company by 2005.

In order of priority:

By 2005, we intend to have at least:

- **25% average annual increase in EPS**
- **25% net income as % of revenues**
- **5 new products/indications approved**
- **5 significant products in late stage clinical trials**
- **\$500 million in new revenues from strategic alliances or acquisitions**

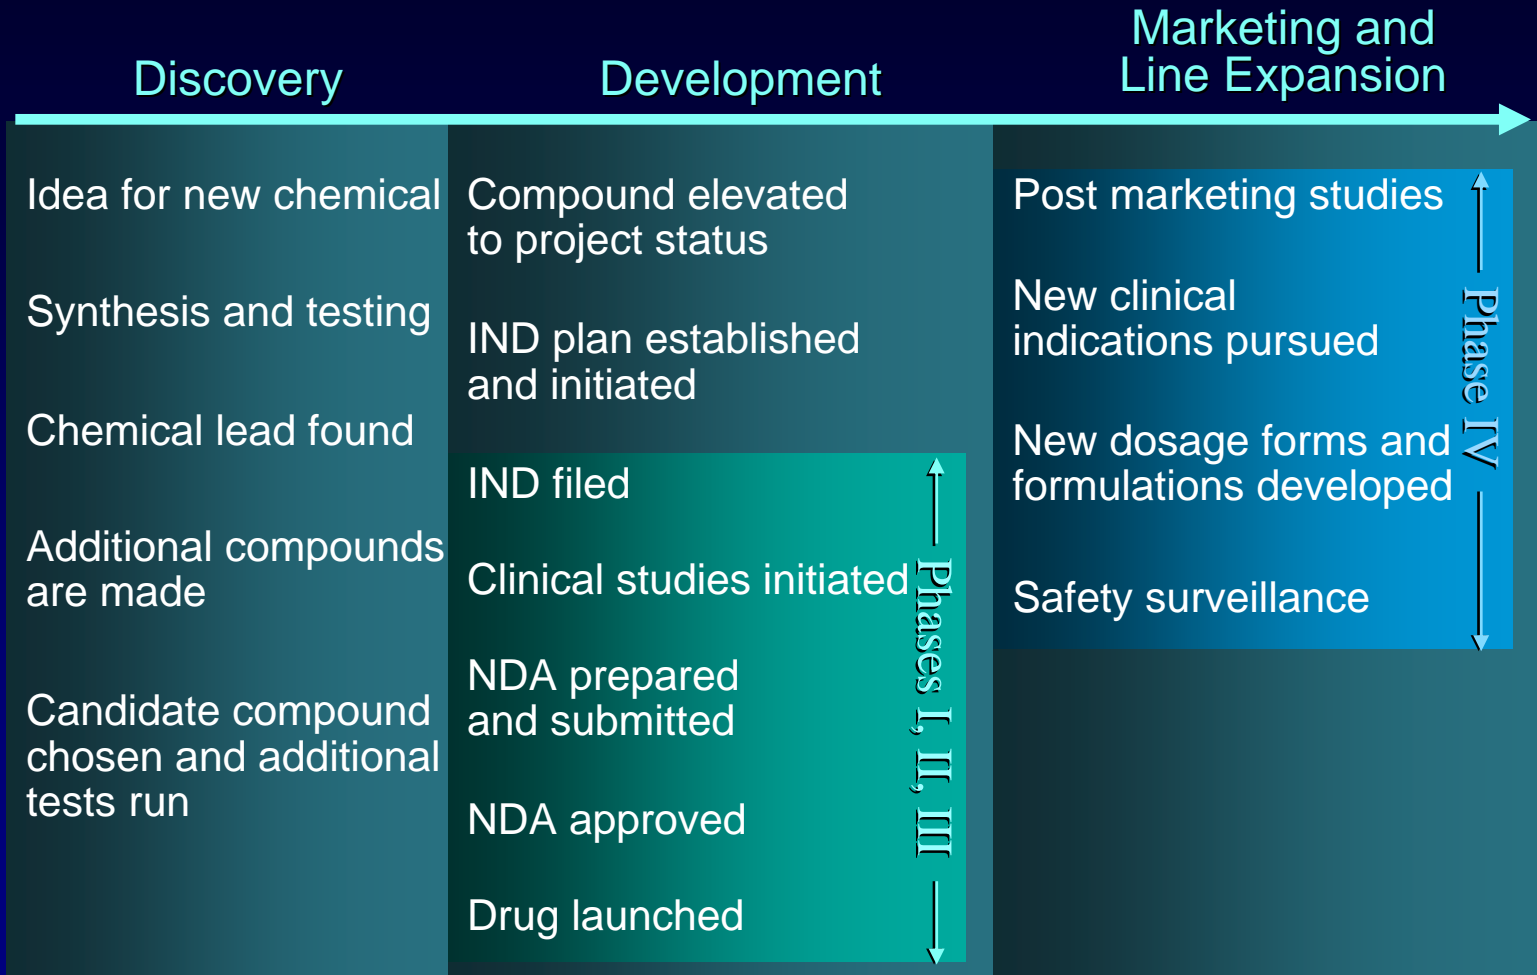
Genentech – Vital Statistics

- ~5,000 Employees



- \$2.2B in Revenue (2001)
 - \$1.9B for first 3 quarters of 2002
- 11 products
 - Protropin[®], Nutropin[®], NutropinAQ[®], NutropinAQ Pen[™], NutropinDepot[®], Cathflo[™] Activase[®], Activase[®], TNKase[™], Pulmozyme[®], Herceptin[®], Rituxan[®]
- 1 product awaiting FDA approval
 - Xolair[™]
- Three major sites
 - South San Francisco, California
 - Vacaville, California
 - Porriño, Spain

Genentech – Process



Genentech – Product Pipeline

Phase I

2C4 Antibody

solid tumors

Anti-Tissue Factor

acute coronary
syndrome

Phase II

MLN-02 Antibody

inflammatory
bowel disease

rhuFab

age-related macular
degeneration

Efalizumab (anti-CD11a)

rheumatoid arthritis

Phase III

Raptiva™

psoriasis

Rituxan®

int/high-grade NHL

Rituxan®

ITP

Herceptin®

adjuvant breast cancer

Avastin™

colon cancer

Nutropin Depot®

adult GH deficiency

Tarceva™


lung cancer

BLA Filed

Xolair™

allergic asthma

Genentech – Life Sciences




...from idea to product

IT Infrastructure

IT Infrastructure




Vacaville



IT Infrastructure

Key

Firewalls & proxies
Web & file sharing
VPN



Supported Platforms

Desktop


- **MacOS 9,X [~2000]**
 - Migrating away from MacOS 9
- **Windows (NT, 2000) [~6000]**
 - Migrating towards Windows 2000
 - Windows XP coming in December
 - NT still in heavy use in validated areas
- **SGI Irix [~40]**
 - Used primarily in Research
 - Molecular modeling
 - Imaging
 - Structural chemistry
- **Sun Solaris [~20]**
 - Used primarily in Research
 - Instrument controllers

Supported Platforms

Servers

- **HP Tru64 UNIX (Alpha) [64]**
 - Infrastructure, Research, Development Sciences, Manufacturing
- **HP HP/UX (PA-RISC) [40]**
 - Finance (Lawson)
 - HR (Peoplesoft)
 - Manufacturing (NovaManage, BPCS, Beckman LIMS)
 - Imaging (FileNET)
- **Linux (IA-32) [34]**
 - Research
- **SGI Irix (MIPS) [4]**
 - Research
- **Sun Solaris (SPARC) [76]**
 - SAS
 - Oracle Clinical
 - Rational
 - Others
- **Windows NT/2000 (IA-32) [300?]**
 - Lots...

Bioinformatics Platform



Protein Engineering / Bioorganic Platforms

Linux Cluster



SGI Origin 2000
(12 processors)



SGI Origin 2000
(16 processors)



SGI Origin 3000
(24 processors)



Supported Databases

Oracle

- Corporate standard
- Used for R&D, Sales and Marketing, Financials, HR, Payroll, IT Services, Facilities, Manufacturing

Informix

- Used for limited legacy applications

SQLServer

- As part of certain applications (black box)

FileMaker

- End-user databases

Access

- Limited use

Custom Databases

- Sequence databases
- Limited applications (BerkeleyDB, MySQL)

Infrastructure Summary

Diversity is the nature of the environment

- Research requires flexibility
- Validation requires more control

Seeing more Linux

- Currently only in Research environments
- Will probably be used in the infrastructure

Significant investment in Tru64 and HP/UX

- Closely watching developments from HP

Oracle is a key component in our infrastructure

Architecture Project

Genentech IT Architecture

Every group had their own needs/ideas

Few points of “control”

- E-mail
- Network
- Firewall
- Central web servers

Lots of technology/vendor favoritism

No process for business unit input into direction

No process for technical consensus

Architecture Project – 1998

Goal:

- Future architecture (3-5 years out)
- Inform future development

Membership:

- Technical staff from around the company

Process:

- 1-day offsite to decide topic areas
- Working groups to propose architectures in each area
- Groups were asked to avoid technology decisions
- 1-day offsite to review results and choose top areas
 - Top areas selected by “vote”

Architecture Project

13 sub-teams:

- Application Architecture
- Automation Architecture
- Database Architecture
- Desktop Architecture
- Disaster Recovery Architecture
- Document Management and Repositories Architecture
- Internet Architecture
- Intranet Architecture
- Network Architecture
- Security Architecture
- Server Architecture
- Software Development Architecture
- Web Architecture

Key Points (Results)

Identified by participants

- Three votes per participant

Top “vote-getters”

- Open systems and standards [11]
- Three-Tier [10]
- Centralized Security, Single-sign on [10]
- Network bandwidth [8]
- Configuration management [8]
- Reusable services and code [6]
- Corporate high availability strategy [5]
- Distributed objects strategy [5]

Key Points (Results)

Top three “vote getters” all result in reduced costs or increased efficiencies:

- **Open systems and standards**
 - Reduction in vendor dependencies
 - Easier integration
 - Quicker new staff integration
- **Centralized Security, Single-sign on**
 - Reduced user time spent dealing with passwords
 - Better security
 - Decreased staffing dedicated to account maintenance in each group
- **Three-Tier**
 - Increased database consolidation
 - Reduced maintenance costs
 - Increased utilization
 - Better uptime & performance

Three-Tier Architecture → Unified Database Tier

Database (3rd) Tier Project

Decided on 2 supported platforms: Tru64 & HP/UX

Started with Tru64 Cluster

Approach

- Single cluster with multiple instances
 - Instances will failover if a node fails
- Instances balanced manually across nodes
- Validated and non-validated databases in separate instances
- Start small, demonstrate stability, performance and value

Implementation

- Built an Oracle-only development/test environment
 - *yukon*: 2-node DS20 Tru64 UNIX Cluster
- Built an Oracle-only production environment
 - *merced*: 2-node ES40 Tru64 UNIX Cluster

“If you build it, they will come”

Offered *merced* as a “third-tier”

- Web applications migrated immediately
 - Were on a single-node Alpha
- Other customers slowly migrated
- Easier than configuring their own, separate third tier

Eventually, became *the* database tier

- Wound up with close to a unified database layer


Architecture allowed for other database servers

- Little interest in additional, separate database servers
 - One group implemented a separate server due to vendor requirements

Current Database Architecture



Oracle Cluster

14 instances
102 databases
> 1TB (growing)



4 instances

Bioinformatics Cluster



Standalone Databases

~10 instances

Web & Application Cluster

Unified Database Tier

Advantages

- Single infrastructure to manage
- Can invest in higher availability
- Separate DB infrastructure can provide more focused tuning
- Significant consolidation

Disadvantages

- Single infrastructure to break
 - All “eggs in one basket”
- Entire platform must be qualified
 - Not bad business practice, anyways
- Upgrades require substantial investment in planning
 - Now on a fixed 6-month upgrade cycle

Challenges

- Customer acceptance
- Version sync
 - different vendors certifying against different Oracle versions

Status

Current databases on unified database tier

- 14 instances across 3 nodes
- 102 databases
 - Finance, Research, Sales, Legal, HR, Product Development, Manufacturing, Marketing, Fermentation, Development Sciences, etc.
- 2 additional databases scheduled to migrate

Databases not on unified database tier

- Medical Affairs (Clinical)
 - Local support group
 - Significant vendor tie-in (Oracle Clinical)
- 5 Research databases
 - Local infrastructure, local support
- 6+ Manufacturing databases
 - Production databases
 - Significant application tie-in
- FileNet, Rational
 - require local databases

Status

Performance

- Completely acceptable
 - I/O Wait times ~0.13 – 0.50 on largest instances

Availability

- Excellent uptime (99.99%)

Most customers are happy

- Still some who want complete control
 - Validation
 - Significant scheduling constraints
 - Comfort

Futures

Integration with Kerberos

- Centralized authentication

Additional databases migrated

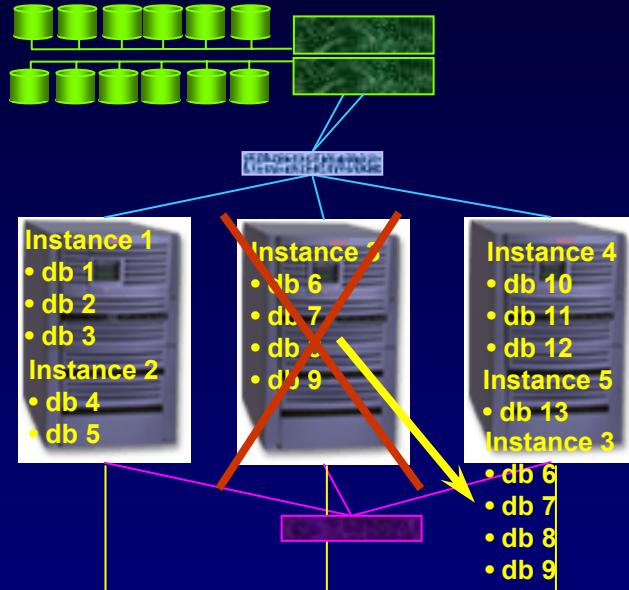
Upgrades

- More nodes
 - If needed
- More disk
 - Will be needed
- Oracle 9i

Oracle 9i RAC

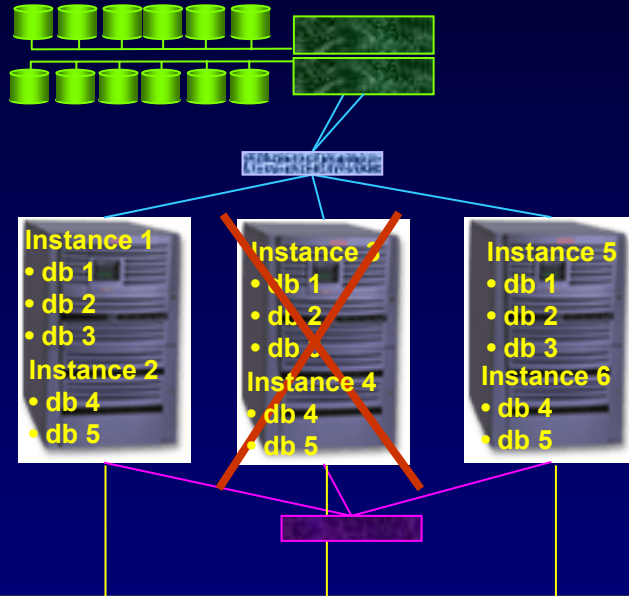
- Oracle 9i RAC Pilot

Oracle 9i



- Manual load balancing
 - Balance by instance
- HA is achieved by failover
 - Failover times 10-30 seconds
- Software must account for disconnects during failover

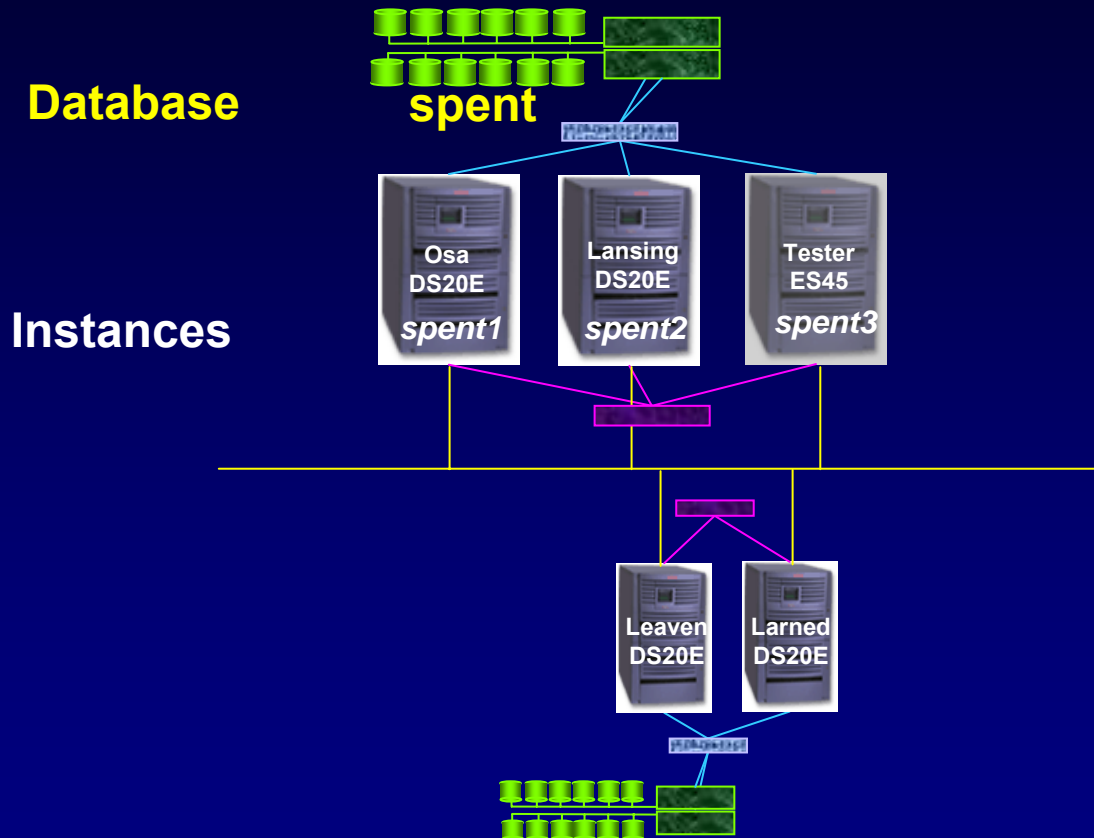
Oracle 9i RAC



- Automatic load balancing
 - Databases served by multiple instances
- Automatic failover
- No software changes to account for node failure¹
- Capacity can be added incrementally

¹Actually, not so sure how invisible this is

Oracle 9i RAC – pilot



Oracle 9i RAC – results

RAC is configured and running on cluster

- 1 database
- 2 instances
- No special tuning (OS or Database)

Load balancing is configured and working

- Least loaded instance gets connected
- Connection-oriented load balancing at this point

Lessons learned:

- It works, but it takes some work
- Start with a demo database
- Getting installation right takes a couple of times

Next steps:

- Ask HP for demo database
- Do detailed performance benchmarking

Oracle 9i – timeline

Too late for current upgrade cycle

Plan to move into production in August

9iRAC

- Technology looks *very* interesting
- Very enthusiastic about using in the future
- Current infrastructure is too good...
 - No business driver to upgrade
 - Current uptime acceptable
 - Current performance acceptable
- Will migrate as business needs dictate

Conclusions

Genentech is all about “*life sciences*”!

IT supports overall corporate mission

IT architectures have evolved

- Decentralized
- More centralized

Unified database tier an important part of our evolution

Questions?

Thank you!

Acknowledgements:

Sonja Bock

Jim Lola


Paul Meadow

Architecture Project

User's View

- Anywhere, anytime computing
- Choice of platform
- Single sign-on
- Integrated corporate applications
 - Only need to enter information once
 - Only need to look one place
- Integrated desktop applications
 - Messaging
 - Conferencing
 - Scheduling
 - Data sharing

Architecture (Systems View)



Architecture

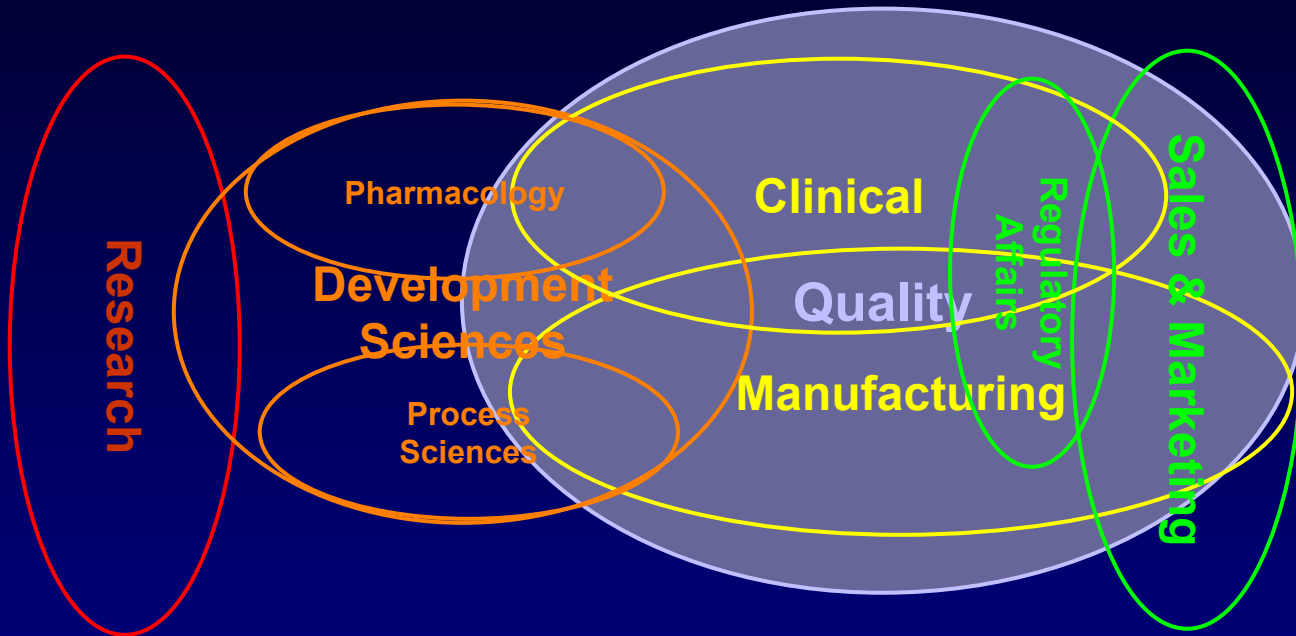
Support View

- Open standards-based
- Centralized where it makes sense
- Decentralized where it makes sense
- Easier deployment and tuning (three-tier)

Developer's View

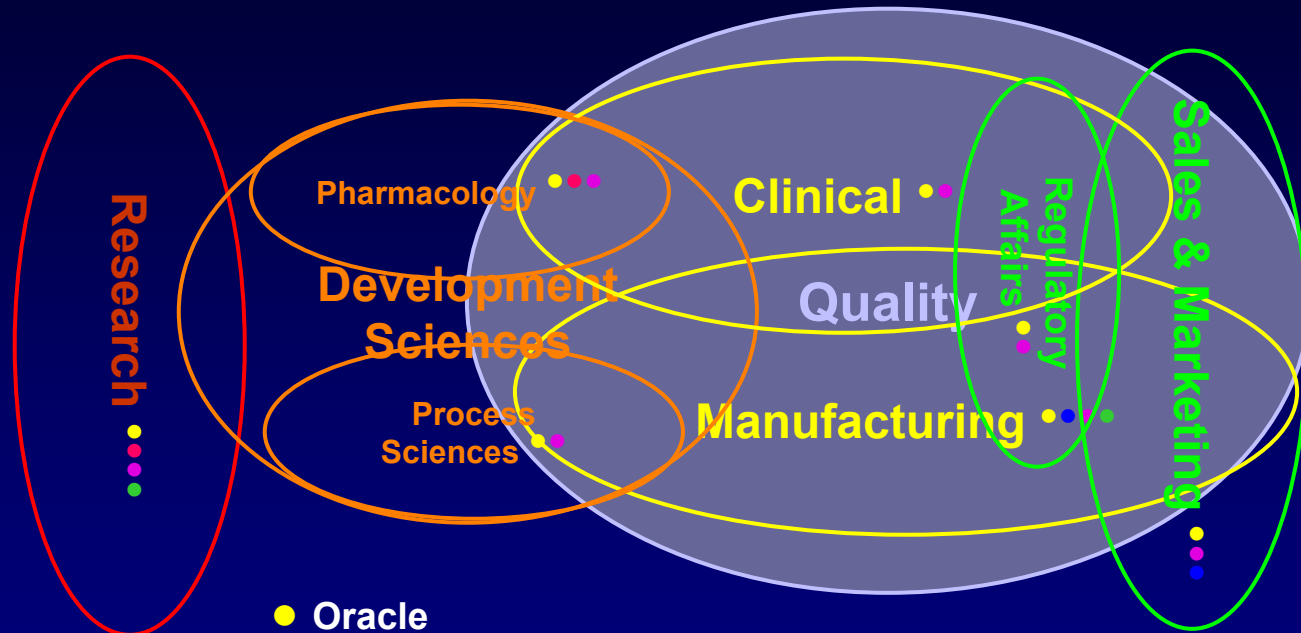
- More code sharing
- Easier integration (distributed objects)
- Better management (configuration management)

Genentech – Pathway



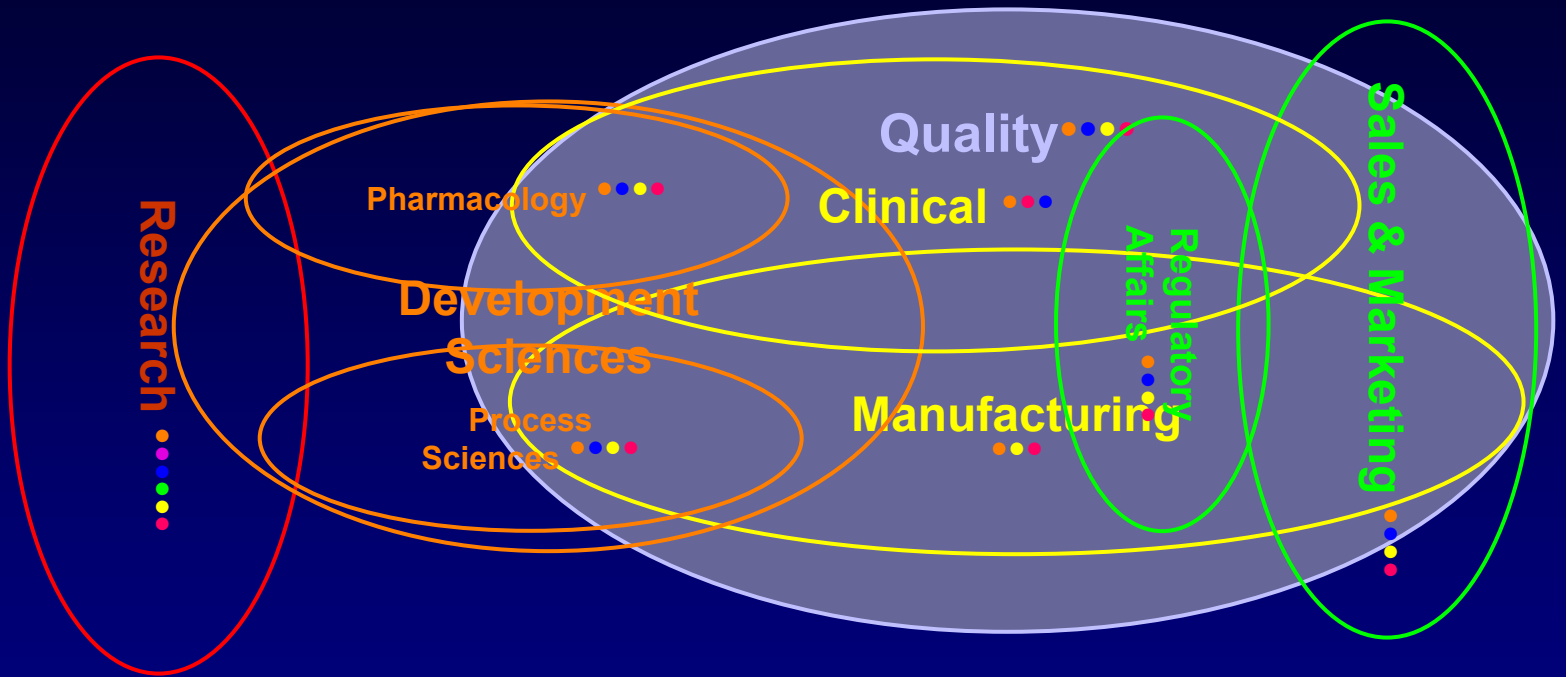
...from idea to product...

Genentech – Databases



- Oracle
- Informix
- SQLServer
- FileMaker
- Custom Databases

Genentech – Platforms



● HP Tru64 UNIX (Alpha)

● HP HP/UX (PA-RISC)

● Linux (IA-32)

● SGI Irix (MIPS)

● Sun Solaris (SPARC)

● Windows NT/2000 (IA-32)