NFS INDUSTRY CONFERENCE

From NFS to NFS

More N, less F, a different kind of S

Rob Gingell

Chief Engineer, Sun Fellow & Vice President

Sun Microsystems, Inc.

gingell@Sun.COM

NFS INDUSTRY CONFERENCE

Outline

- Classic NFS
- Evolution:
  - Networking
  - Applications (Files)
  - Systems
- One company's response
- Collective Challenge

NFS INDUSTRY CONFERENCE

Classic NFS

- Client-server network application
  - A factoring of “OS” functionality
  - Decomposition and distribution
  - Allowed sharing, required heterogeneity
- Associated system services
  - NIS – personality of “UNIX”
  - Judicious: network aware
    - But not overly presumptive
    - *NOT* a distributed OS: contrarian in mid-1980's

NFS INDUSTRY CONFERENCE

Classic NFS (cont.)

- Satisfactory Evolution
  - Wideband support
  - Broadband support (DAFS/RDMA)
  - Heterogeniety: attributes
  - Authentication, sort of
- Suitable for organizations
  - Even very large ones

NFS INDUSTRY CONFERENCE

Network Evolution

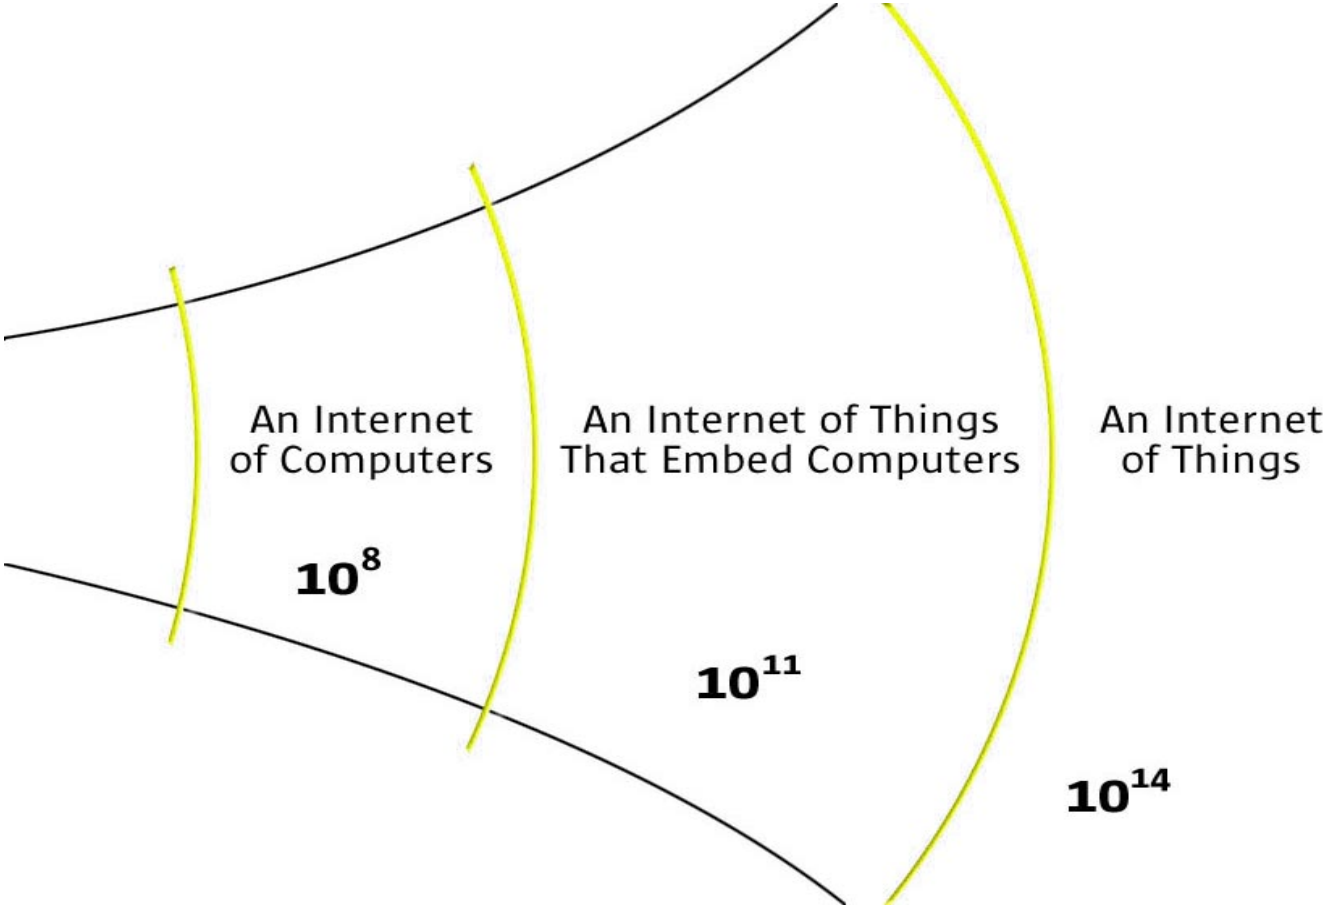

NFS INDUSTRY CONFERENCE

Network Scaling

- Bandwidth/Cost Improvements
  - Changes engineering assumptions
  - Allows rule-changing when costs change
- High Expectations
  - Never rebooted
  - Continuous application operation
  - But, something is always broken
- Exposes flaws in thinking

NFS INDUSTRY CONFERENCE

Deutsch: Fallacies of Networking

- The network is reliable
- Latency is zero
- Bandwidth is infinite
- The network is secure
- Topology does not change
- There is one administrator
- Transport cost is zero
- The network is homogeneous

NFS INDUSTRY CONFERENCE

NFS Fallacies

- " NFS server not responding..."
- AUTH_UNIX
- I made my application more robust
  - How come it's running so slow now?
- Automounting
  - Every problem solved by an indirection
- Transparency (of the network) considered harmful

NFS INDUSTRY CONFERENCE

Application Evolution

Client/Server

Objects

Web Application

Web Service

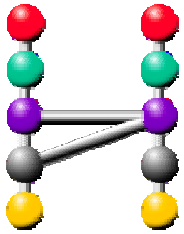

Next

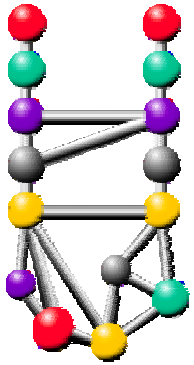

Next Next

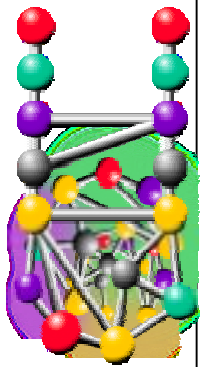

NFS INDUSTRY CONFERENCE

Application Evolution (2)

- " Style" of applications evolving
  - Transfer to transaction
  - Now content, soon telemetry and control
  - It's not going to stop: change is constant
- Applications more acquisitive
  - " A photon is a terrible thing to waste..."
  - Operating on objects, not files
  - Streams, media, and yes, files too
  - Need pumps, 3rd party motion

NFS INDUSTRY CONFERENCE

Application Evolution (3)

- Arriving as a “system”
  - Distinct level 5/6 protocols
  - Carrying own notions of personality
- Security: Living on the internet
  - Self-contained perimeters, e.g., VPN
- Transparency: semantics needed
  - Applications participate in availability
  - Transactions a terrible thing to lose

System Evolution

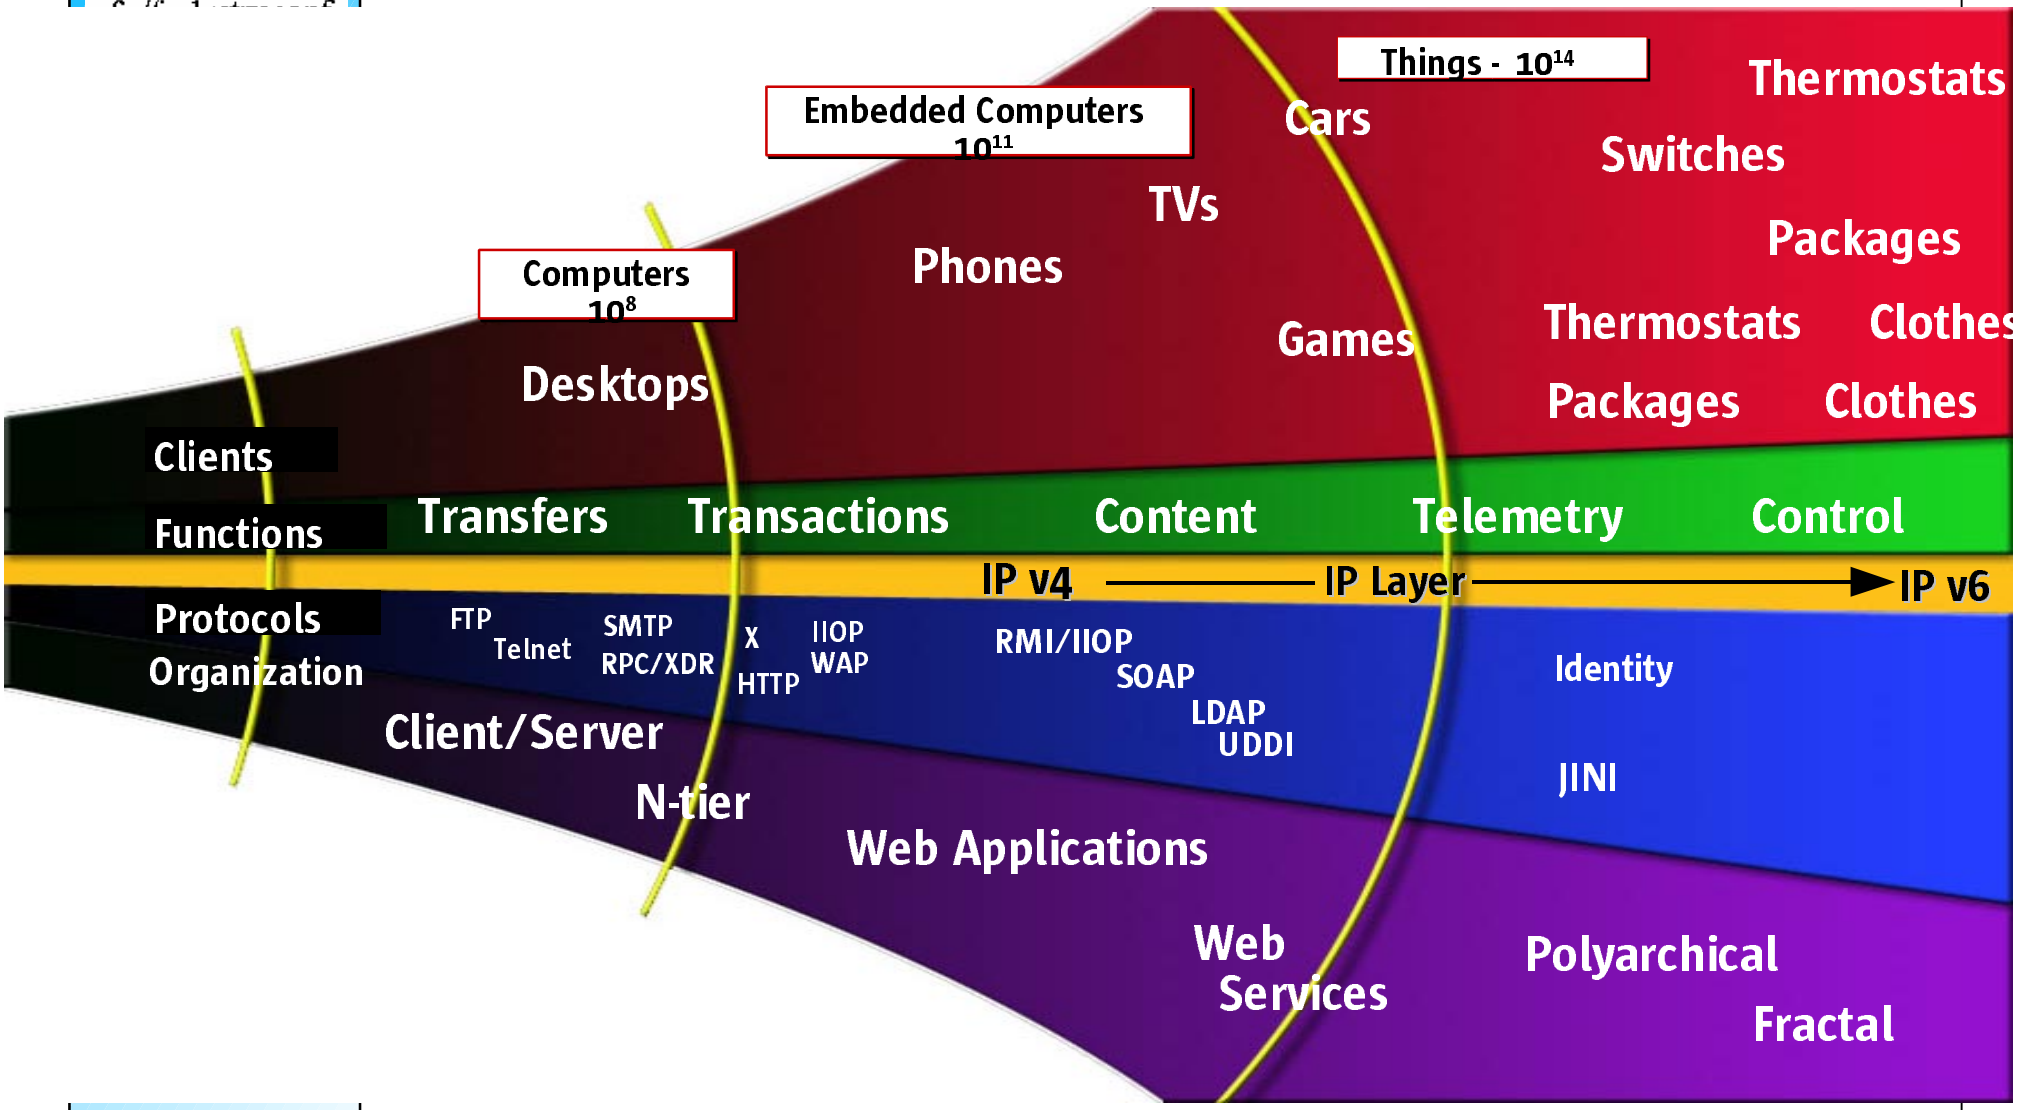

NFS INDUSTRY CONFERENCE

Organizing Principles

- OS is not an “organizing system”
  - Access devices
  - Application systems transcendent
  - Administrators: of the environment
  - IP-based protocols provide organization
- Flexibility
  - Virtualization of resources, provisioning
- Security
  - Not trusted base, trust relationships

NFS INDUSTRY CONFERENCE

Organizing Principles (2)

- Transparency: no
- Complexity mitigation: yes
- Factor, not weld or “layer”
  - NFS was not: **`if (network)`**
  - NFS was: VFS, generalization
- Scale: distribution/dispersion/size
  - Implies: lots of copies to achieve virtualization, reliability, availability
- Technology: rust is really cheap

NFS INDUSTRY CONFERENCE

Sun's Issues/Needs

- Network is evolving rapidly
  - Not like boom years, but still going
- Application shift is occurring
  - Developer pool order of magnitude larger
  - Primarily (but not exclusively) Java
- Network inherent benefits
  - Flexibility, robustness, but “sophistication”
- Diversity (of use, things) increasing

NFS INDUSTRY CONFERENCE

Sun's Issues/Needs (2)

- Current systems compose poorly
  - Inefficient for network applications
  - Measured by service/assembly cost
- Therefore, we want N1:
  - Systems that deploy more naturally
  - Network applications as design center
  - Elements linked by IP protocols
  - Storage environment by definition: “NFS”

NFS INDUSTRY CONFERENCE

Not your father's NFS

- More than files, less than files
  - A factoring opportunity
  - Segregating: naming, pumping, caching
  - Aesthetic: stop at greatest common factor
    - But expect ongoing iteration
  - A toolkit for application “file systems”
- Personal Example: **`mmap()`**
  - “ Less” than **`read()/write()`**
  - But, a building block for them

NFS INDUSTRY CONFERENCE

Not your mother's NFS either

- Confronts fallacies by design
- Reliability: no single point of failure
- Performance: policy based
- Transparency: QoS and policy
  - For administrators, not programs
- Security: another opportunity
- Flexibility: a network advantage

NFS INDUSTRY CONFERENCE

Is it actually NFS?

- Probably not just NFSv$_{next}$
  - At least, not just a • to the protocol
- More a set of building blocks
  - Reusable concepts, GCF's
- But: occupies same concept space
  - Network Storage System more accurate
  - But higher, not lower, concepts
    - IP-SCSI, blocks are *so* not it
  - NAS box evolution

NFS INDUSTRY CONFERENCE

What's the Specification?

- I don't know: starting journey
  - Journey must be taken
  - Know what it must do to end
- How will it happen?
  - By taking a cut at making a real system
  - Part of a larger overall systems program
  - Focus: solving storage problem
- Where will it happen?
  - Maybe Sun and/or not/also, up to you

NFS INDUSTRY CONFERENCE

Collective Challenge

- NFS: legacy of the OS world
  - Network attached, not network organized
  - Very successful
    - Most successful client/server application?
  - Isn't going away
- Need: post client-server “NFS”
  - Generalization of NAS
  - Applications trend: *very* storage intense
  - Opportunity is tremendous