



PRESENTS

NETWORLD INTEROP



The New WAN Link: Dynamic Information and Hierarchical Caching Solutions

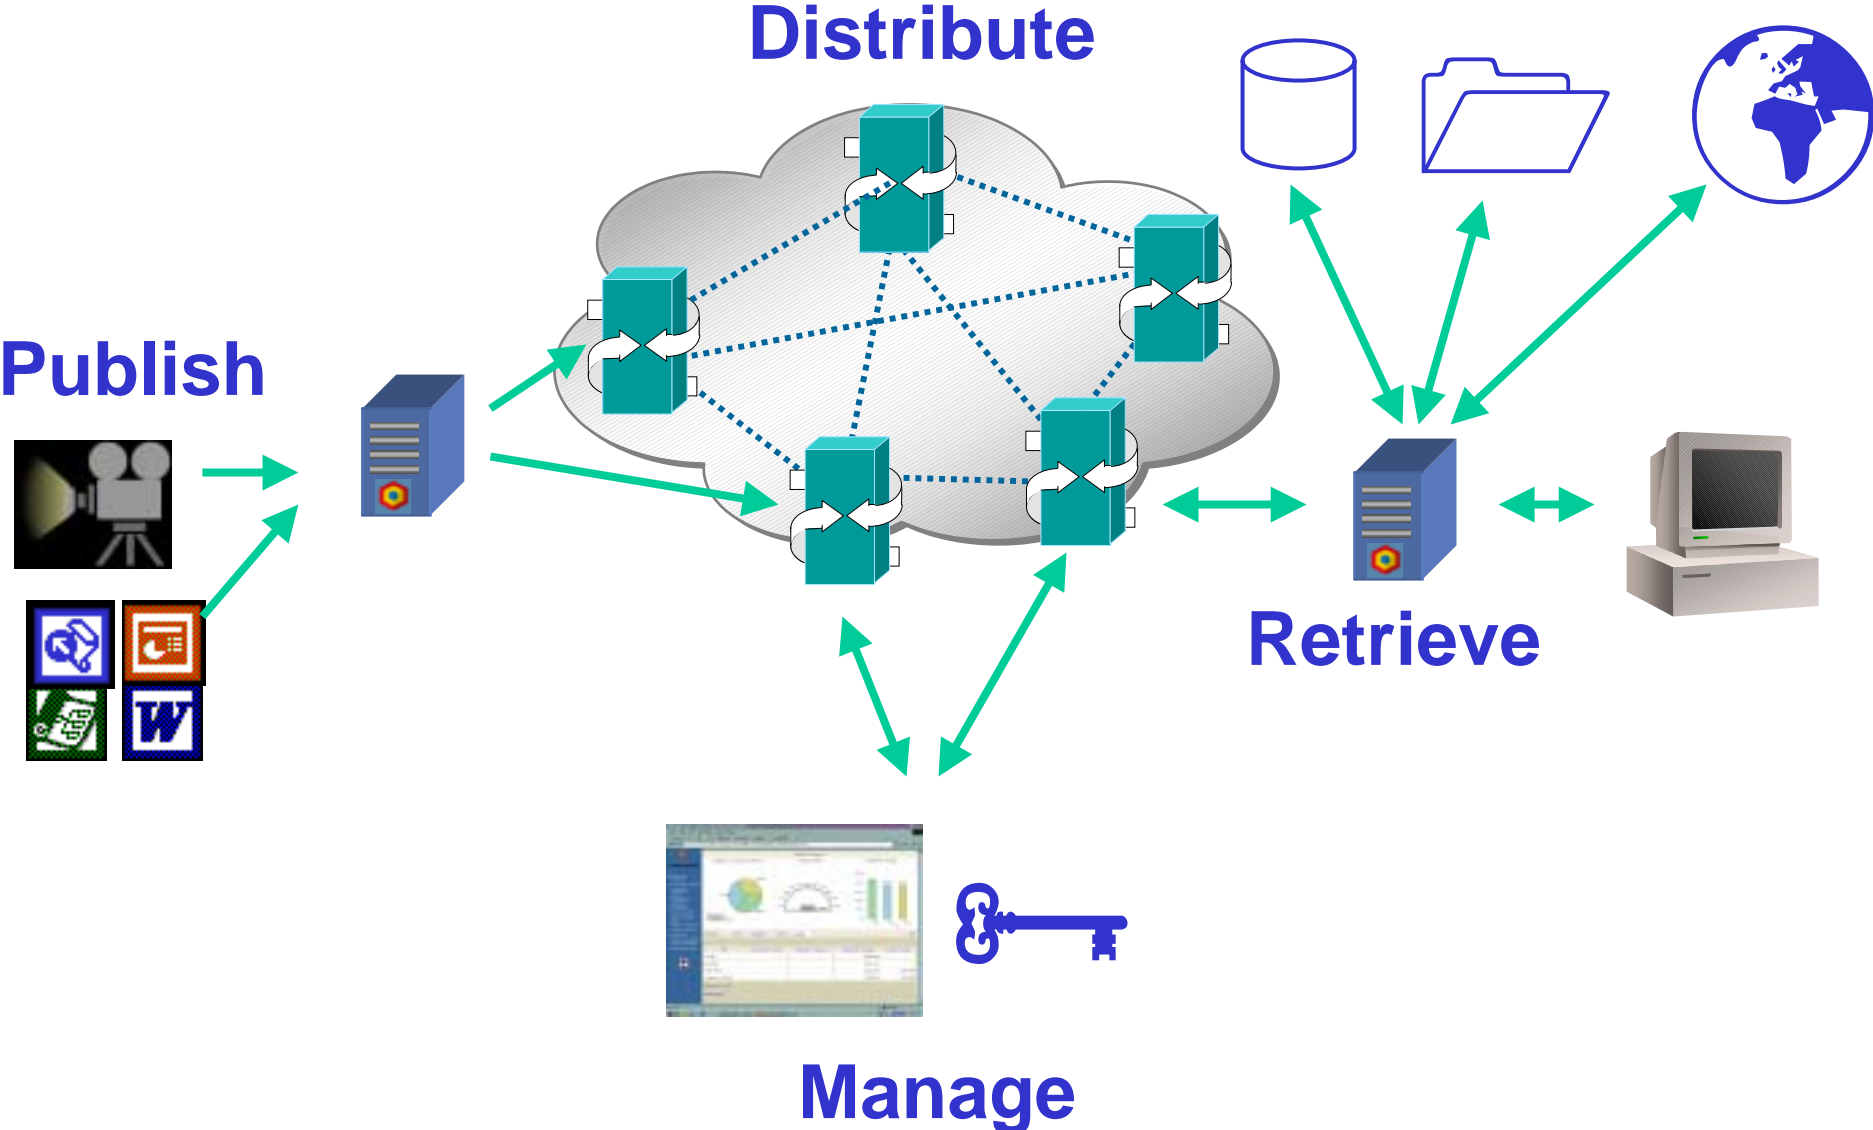
Keith Weng

Senior Strategist, Inktomi

12 Sep 2001



End-to-End Architecture



Background and Experience

- Senior Strategist, Inktomi
- Founder of Webspective 1997:
 - Content distribution & web infrastructure management
 - Acquired by Inktomi 1999
- Founder of Open Environment 1992:
 - 3-tier client/server -- DCE, Corba
 - IPO'd 1995, acquired by Borland
- Product design; systems design and onsite implementation with customers; systems administration, release management
- Technical consulting on client/server and Web infrastructure to Fortune 500: Motorola, Fidelity, American Express, BankOne, Compaq, Enron, John Deere, Sprint, Chrysler, Digital Island/C&W ...

Real-Life Deployment Required for N-Tier Computing

1991: 2-tier client/server, Powerbuilder, VB, oracle, Sybase, over dialup and on LAN, **TCP/IP** stacks, win3.0, LanManager

1994: 3-tier, DCE, **CORBA**, DCOM, over dialup, LAN and WAN, win3.1

1995: public internet, corporate **email over internet**, software **firewalls**

1996: cgi, java, applets, **NTLM**, win95, DHCP, (Corba/DCOM can't pass firewalls)

1997: **SSL** for credit cards, javascript, T-1's, hardware firewalls, Palm, **kerberos**

1998: **Transparent caching**, 100Mbps ethernet, **LAN load balancers**, LAN switches, **LDAP**

1999: **app servers**, EJB's, VPNs, T-3's, **DNS load balancers**, geographic data centers, (but db's can't replicate over WAN)

2000: **CDNs**, b2b automation, WAP, scalable app servers,

The Dynamic Caching Dream

“CDN for Dynamic Content”

- All applications Webified, in particular mission critical ones
- Closest cache to user handles anything redundant (including subcomponents of dynamic and even personalized pages)
- Application session failover and redundancy supported
- Security still assured (all of AAA)
- Origin servers / data centers don't have to scale up
- Thin pipes spared
- User experience improves
- Outsource scaling of applications to CDN's
- SLA's on whole service

The Dynamic Caching Reality

- Some “page assembly” technology available for LAN use, some internet services available for non-sensitive
- But:
 - Web Applications must be modified for page assembly
 - At minimum tagged, often rewritten or restructured
 - More complex to test than static caching/CDN, standalone app.
 - Security ACL granularity doesn’t easily correspond for caching
 - No delegated security across administrative domains
 - Existing standards efforts don’t form a solution (XSLT/XSL/CSS, RUP, ESI, ...)
 - No management tools for many-WAN-node dynamic content
- So today usable for least sensitive, least critical content
- How to solve for more sensitive, more critical ? ...

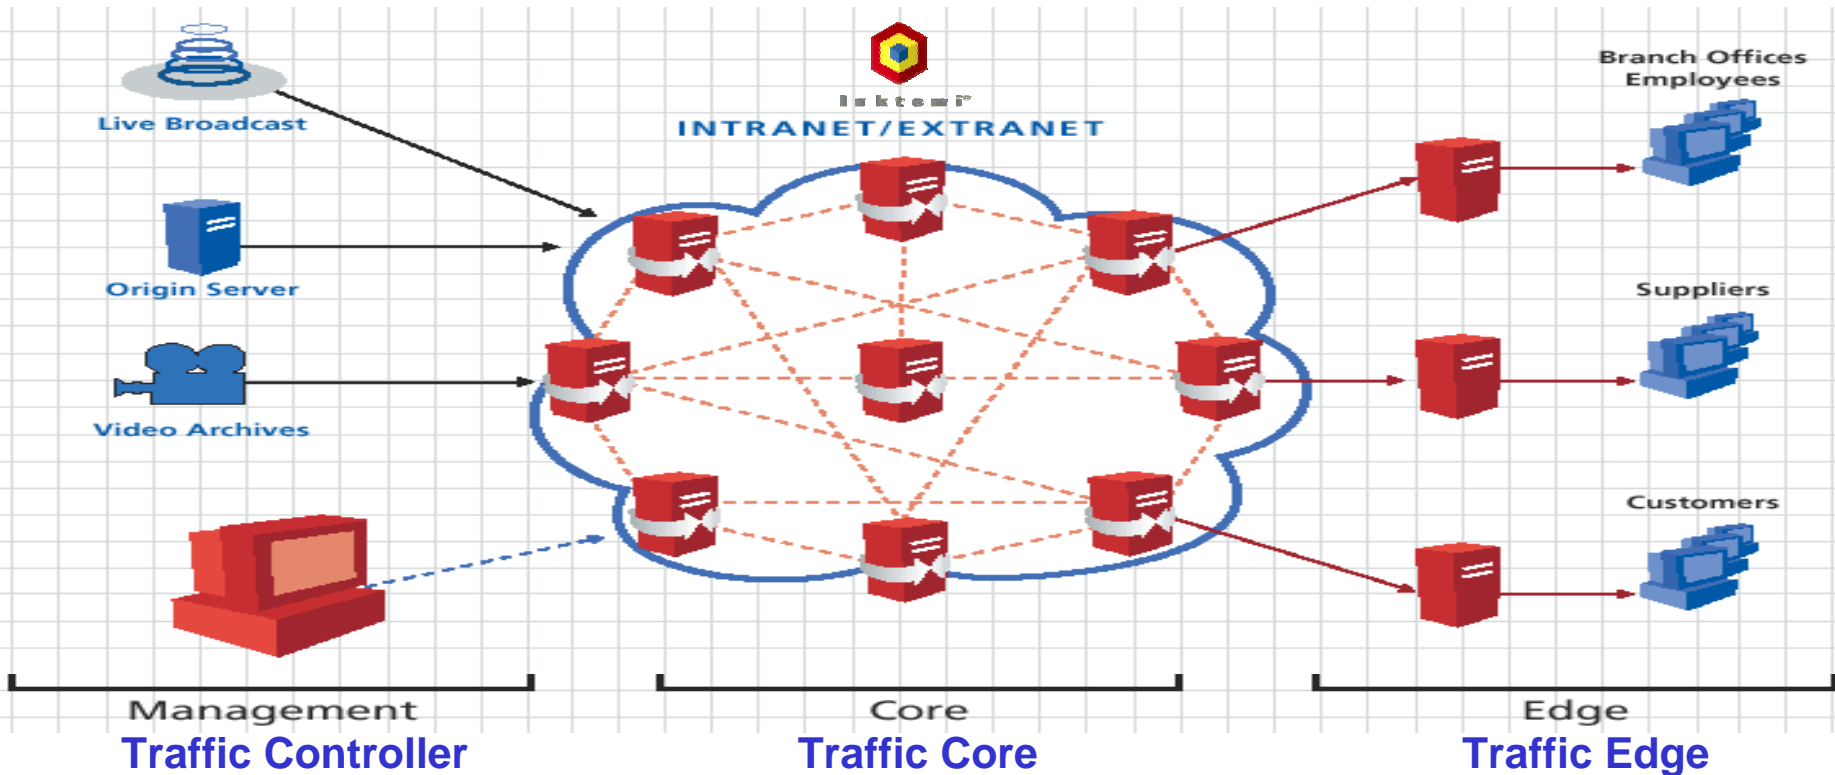
Biggest Issues

- Control and trust across administrative domains
 - b2c: Fidelity stock trading -> Sprint -> Earthlink -> user device
 - b2b: GM procurement -> Uunet -> Genuity -> NapaAutoparts
 - How can app owner **trust** what happens in network twice removed?
 - E.g. Guarantee accurate access controls on sensitive data on a cache in Earthlink POP or in NapaAutoparts datacenter
 - How can app owner **control** what happens there ?
 - Eg. Control when cache invalidations or preloads are to be cascaded to all hierarchical internal caches within NapaAutoparts ?
- Scalable, reliable coordination across many WAN processing nodes
 - How to update content in an event-driven way (not just expiration)?
 - How to manage lights-out closet boxes for *application* processing, not just caching?

Possible Answers

- Across administrative domains (extranet and internet):
 - Proprietary CDN with API's
 - Have ESI now, but when with AAA? Lockin?
 - Standards for remote page assembly and AAA (slow to develop)
 - Need multi-network SLA's on AAA? (never?)
 - .Net/hailstorm
 - Everything a pay-for-use service, everyone meet in the middle
 - Actually probably still need caching at edges
 - Build it yourself with middleware and caching components
 - Cost, reliability?
 - End-to-end solutions

Example: End-to-End Solution for Streaming



- Monitor all nodes
- Control content updates
- Aggregated logs
- Control QoS

- “Application-Level Networking”
- True fail-around, not splitters
- Manage 000’s of subs
- Manage 000’s updates
- QoS, prioritize management

- Cache multiple formats
- LDAP integration
- Feedback mgmt stats
- Remote upgrade