



Rise of the Appliances

How purpose built hardware is enabling the next generation of massive scale infrastructure

Hans Jespersen
Principal Systems Engineer
Solace Systems Inc.

hans.jespersen@solacesystems.com

Simple building blocks...



...taken to the extreme!



Lego house built for Top Gear's James May

The (not so) Central Processing Unit




HD H.264 Video CODEC



Remember when video encoding was a batch process?


There's much more to Moore than MIPS



Let's take a techie tour



FPGA



Xilinx Virtex-7 2000T FPGA

- More than twice the capacity and bandwidth offered by the largest monolithic devices
- 2 million logic cells (equivalent to 20 million ASIC gates)
- 6.8 billion transistors

Verilog Hardware Description Language

Example

[\[edit\]](#)

A hello world program looks like this:

```
module main;
  initial
  begin
    $display("Hello world!");
    $finish;
  end
endmodule
```

A simple example of two flip-flops follows:

```
module toplevel (clock, reset);
  input clock;
  input reset;

  reg flop1;
  reg flop2;

  always @ (posedge reset or posedge clock)
  if (reset)
  begin
    flop1 <= 0;
    flop2 <= 1;
  end
  else
  begin
    flop1 <= flop2;
    flop2 <= flop1;
  end
endmodule
```

Horizontal scalability on a single chip




Intel Westmere-EX

- 2.6 billion transistors
- 10 64-bit cores @ 2.4 GHz
- 7,200 MIPS
- 130 watts TDP

Xilinx Virtex-7 2000T FPGA

- 6.8 billion transistors
- 3,600 8-bit processors @ 100 MHz
- 180,000 MIPS
- 20 watts TDP

GPU



\$239.99 = 275,000,000 polygons per second



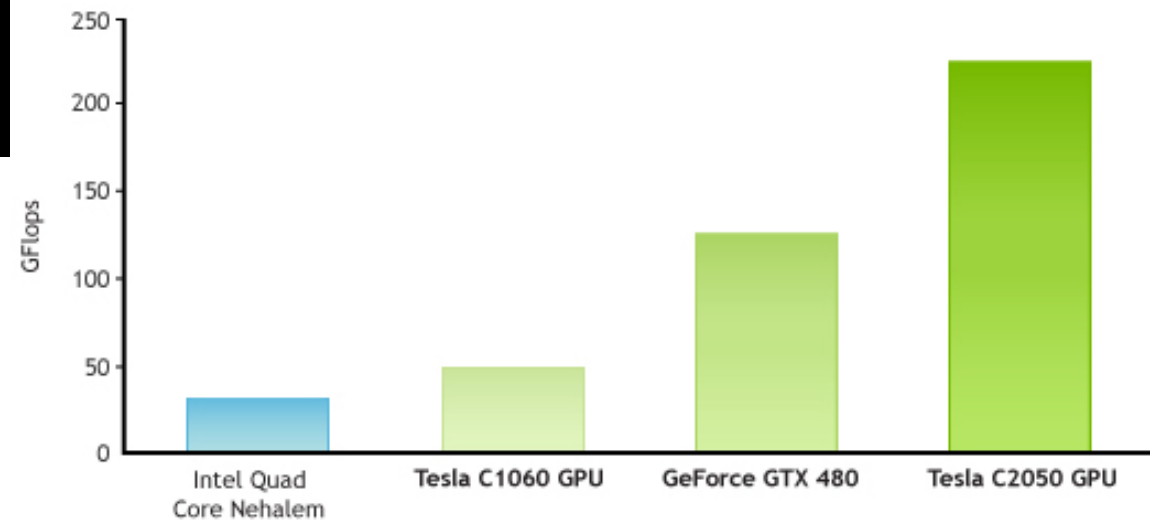
Supercomputing Floating Point Performance



Tesla M2090

665 Gigaflops double precision
1331 Gigaflops single precision
512 Cores
6 GigaBytes Local Memory
177 GBytes/sec Memory bandwidth

N-Body, Double Precision, n= 107,520



Performance of N-body on x86 CPUs and different GPUs

Scaling Interconnects



Mellanox SwitchX

Integration of 144 SerDes (Serializer/Deserializer)

InfiniBand, Ethernet, Fiber Channel and PCIe

- Up to 4Tb/s switching capacity

Flexible Port Configurations

- *Up to 36 56Gig IB Ports*
- *Up to 36 40GigE Ports*
- *Up to 64 10GigE Ports*
- *Up to 24 2/4/8Gig FC Ports*

Imagine:

IB for RDMA between cluster nodes for distributed state

40 GigE for Data Center Ethernet Backbone

1/10 GigE for Internet Facing connections and server farms

Fiber Channel for RAM SAN Storage

Network Processors



Offloading the IP Stack to Hardware



Cavium Octeon II

- 32 core MIPS64 Processor

Pre-built application acceleration engines

- Packet Processing
- Encryption/Decryption
- Deep Packet Inspection (RegEx)
- Compression/decompression
- De-duplication
- RAID

Millions of concurrent connections

The GeoWeb and the GeoRouter


The Geospatial Web or *GeoWeb* is a term used to describe the convergence of Web technologies with geographical (location-based) information.



Location is the fuel driving mobile app growth



What Powers the Internet Growth Rate?



Federated GIS is just in its infancy



Dude, where's my hospital?



The Story of a Simple Mission



“Protect Our Nation from Dangerous People”

Try exchanging “alert” across the following

- **State of New York**
- **State of New Jersey**
- **State of Connecticut**
- **NYC Government**
- **New York City Police Department (NYPD), NJPD...**
- **New York City Fire Department (NYFD), NJFP...**
- **Port Authority of New York/New Jersey (PANYNJ)**
- **Metropolitan Transportation Authority (MTA)**
- **Various Federal Agencies, NGO’s (i.e. Red Cross), etc.**

- **and that’s just the New York area!**

Step 1: Connect Everything



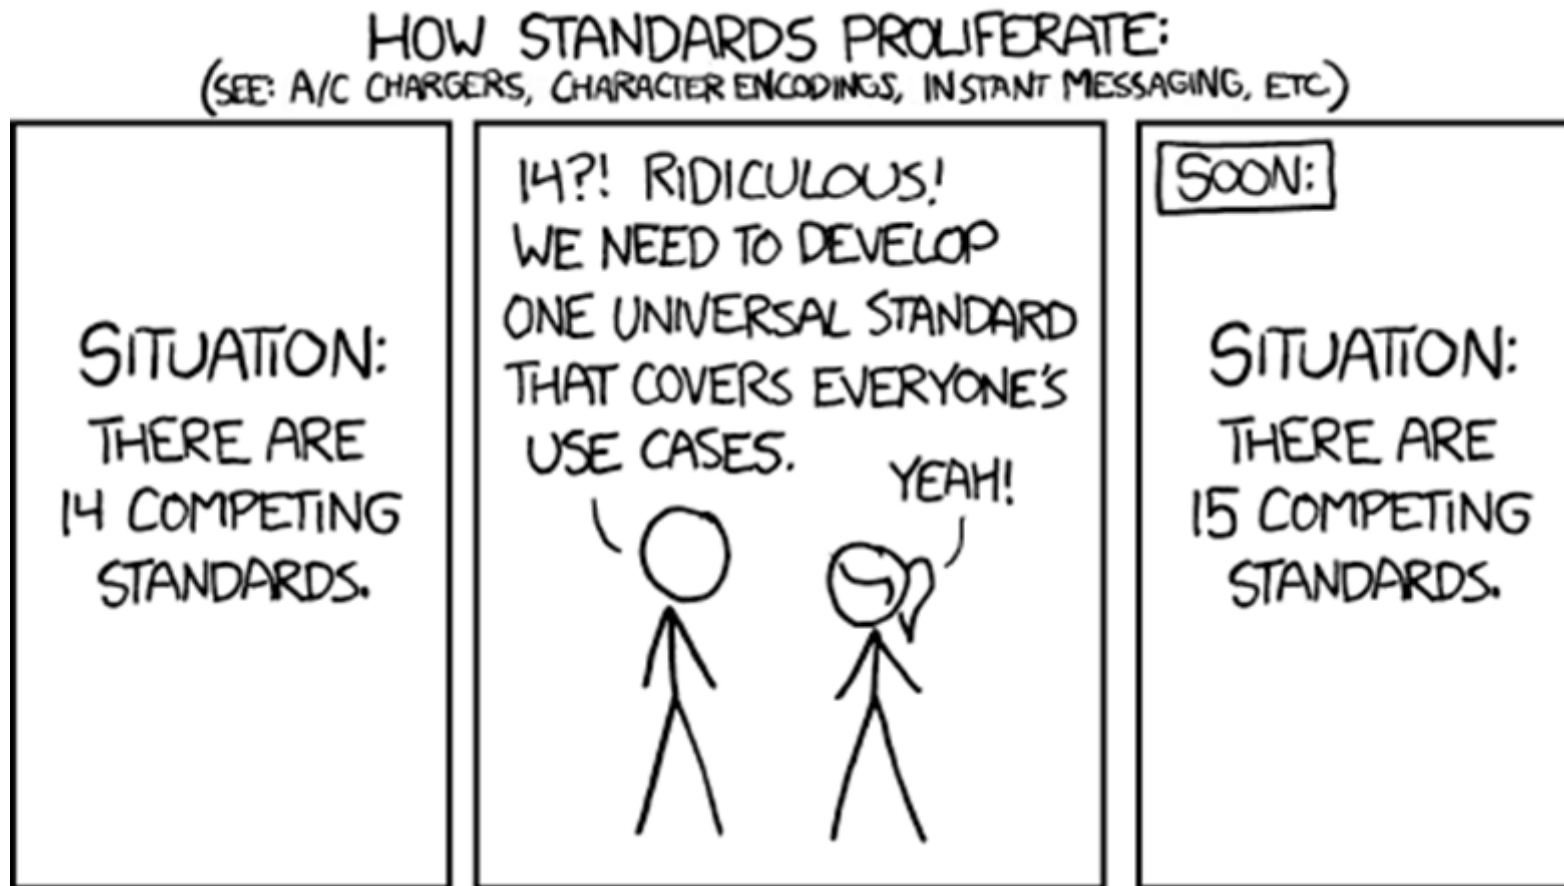
Step 2: Agree on names for everything

© Original Artist
Reproduction rights obtainable from
www.CartoonStock.com



Step 3: Follow all applicable standards

- *CAP v1.1, CAP v1.2, CAP with IPAWS Profile, CAP-CP, EDXL-DE, EDXL-RM, EDXL-HAVE, EDXL-SitRep, Cursor On Target (COT), NIEM CBRN N.25, NIEM SAR, ANSI N42.42, IEEE 1512.1, IEEE 1512.2, IEEE 1512.3, UN/LOCODE 2007, ...*



<http://xkcd.com/927/>

Federated ESB : The Problem



The Power of the Point and the Polygon



Proximity


*Within a half mile of
The White House*

Point

*The Washington
Monument*



Polygon



We can make lots of cool things with polygons



DHS/DNDO Mission

The screenshot displays a Google Earth interface with a mission route plotted along the East River. The route is a thick, multi-colored line connecting seven points labeled A through G. The map shows the East River, the East River Bikeway, and surrounding urban areas with street names like Cherry St, South St, and Front St. The Google Earth interface includes a search bar, a 'Places' list, and a 'Layers' panel.

Search Bar: Fly To: Find Businesses, Directions. Fly to e.g., 37 25.818' N, 122 05.36' W.

Places List:

- QASIS Demo
- Federal Demo
- Manhattan Island
- Central Park
- Bounding Box
- distribution/range
- A 40°42'20.41"N 73°59'46.64"W
- B 40°42'30.92"N 74° 0'0.04"W
- C 40°42'12.92"N 73°59'37.26...
- D 40°42'22.67"N 73°59'53.05...
- E 40°42'33.53"N 73°59'43.20...
- F 40°42'6.69"N 73°59'44.95"W
- G 40°42'23.93"N 73°59'24.81...

Layers Panel:

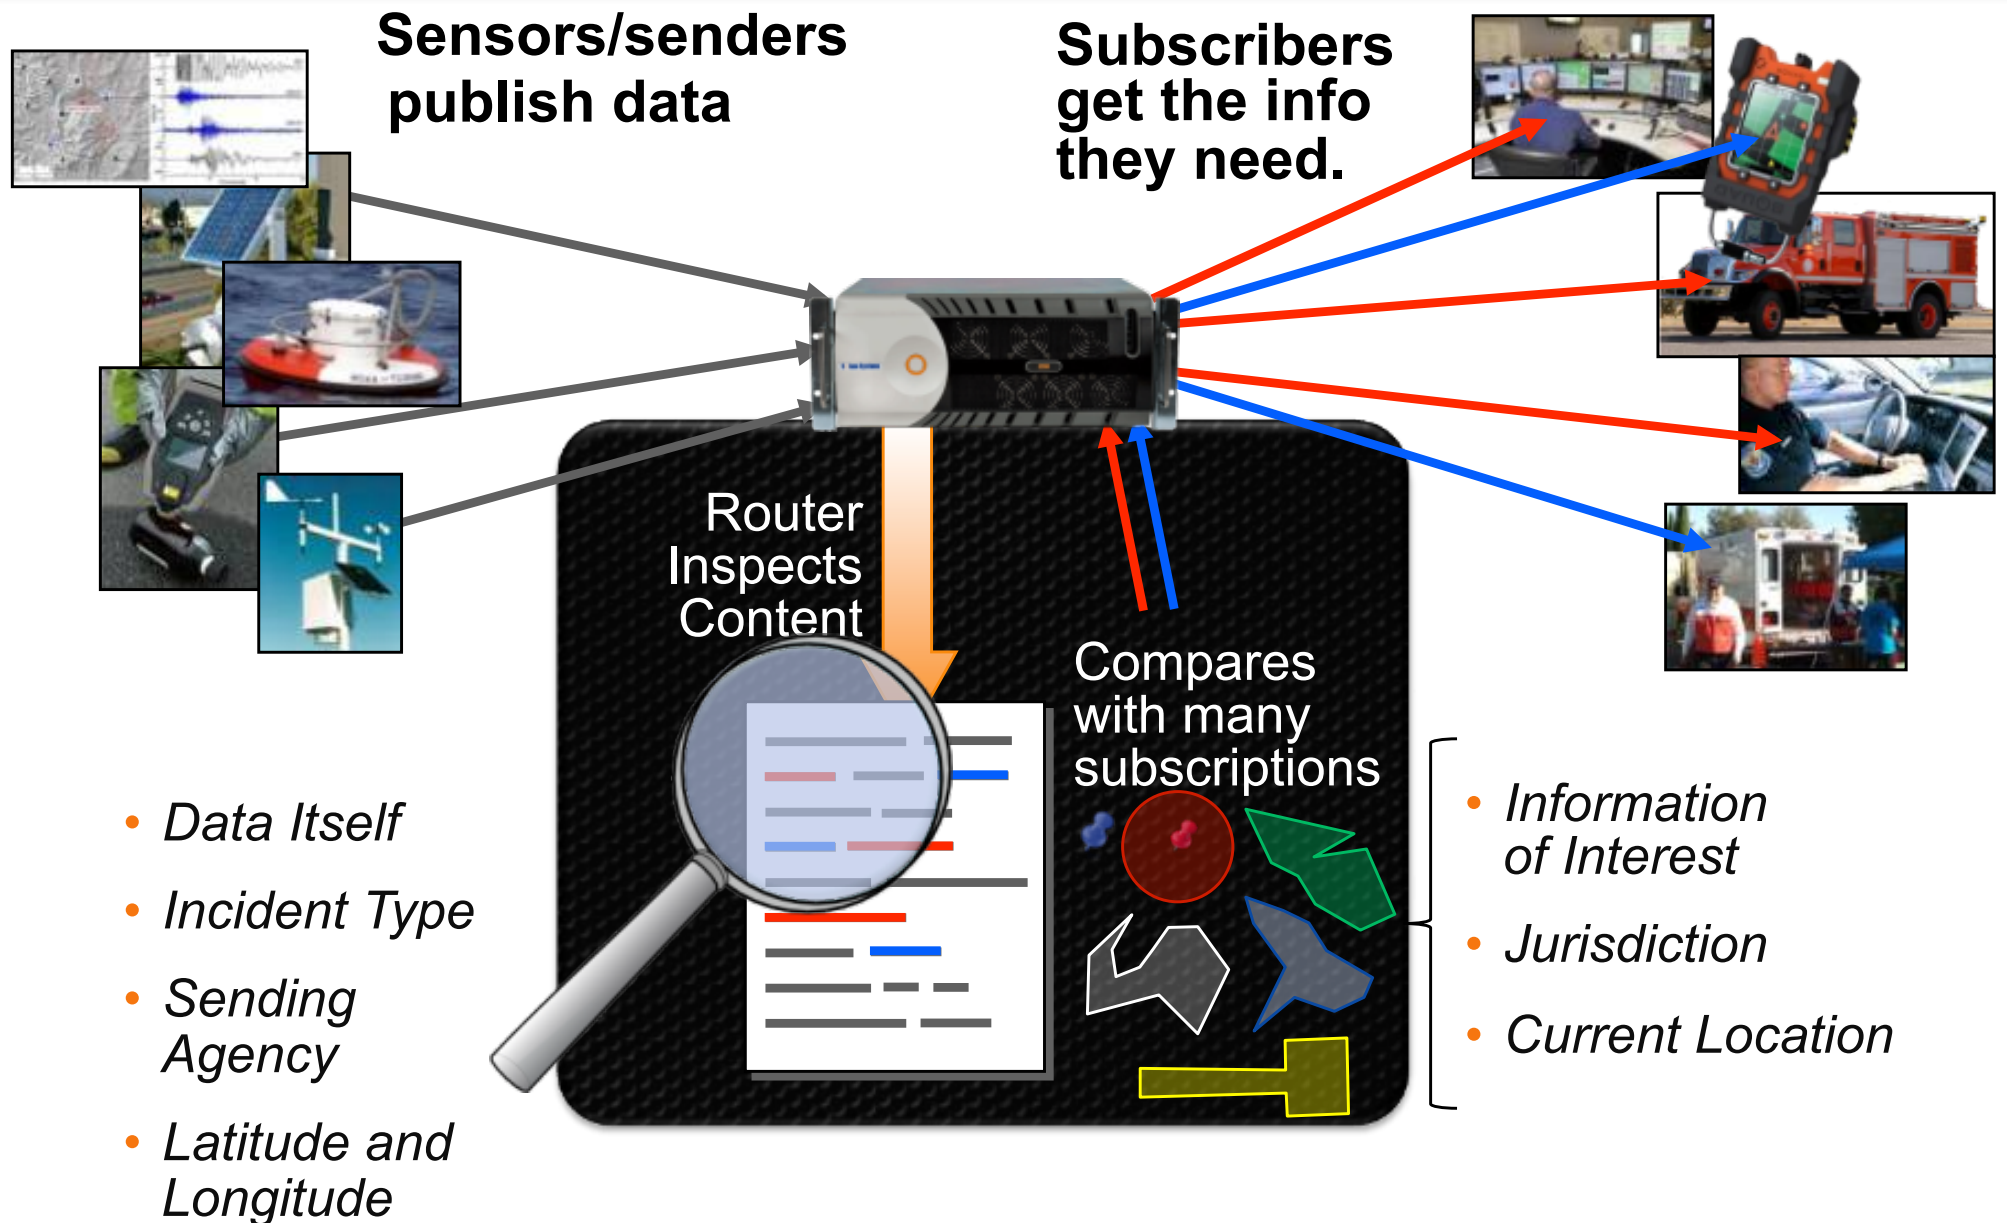
- Primary Database
- Geographic Web
- Roads
- 3D Buildings
- Street View
- Borders and Labels
- Traffic
- Weather
- Gallery
- Ocean
- Global Awareness
- Places of Interest
- More
- Terrain

Map Coordinates:

- Point A: 40°42'20.41"N 73°59'46.64"W
- Point B: 40°42'30.92"N 74° 0'0.04"W
- Point C: 40°42'12.92"N 73°59'37.26"W
- Point D: 40°42'22.67"N 73°59'53.05"W
- Point E: 40°42'33.53"N 73°59'43.20"W
- Point F: 40°42'6.69"N 73°59'44.95"W
- Point G: 40°42'23.93"N 73°59'24.81"W

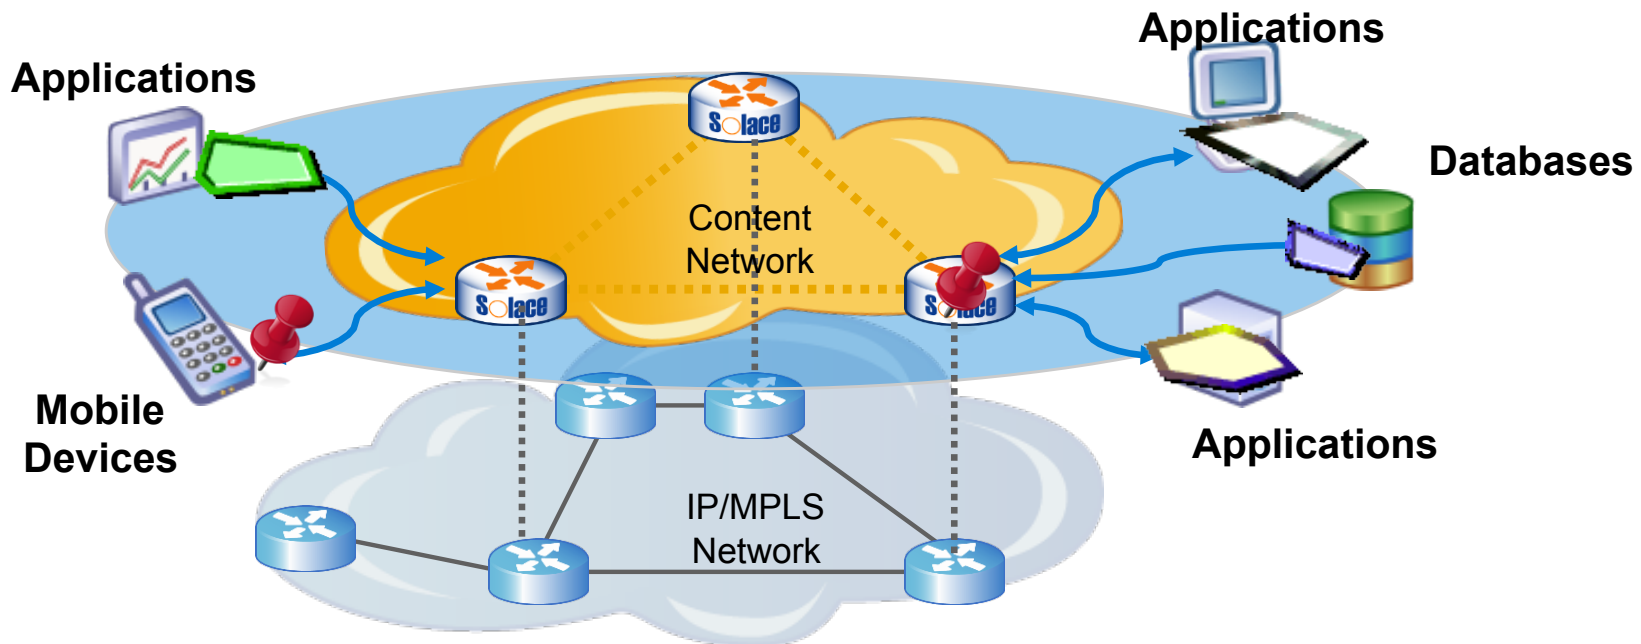
Map Footer: Image © 2010 Sanborn, Image © 2010 DigitalGlobe, © 2010 Google. Eye alt 1.23 km.

Real-time Geospatial Routing in Hardware

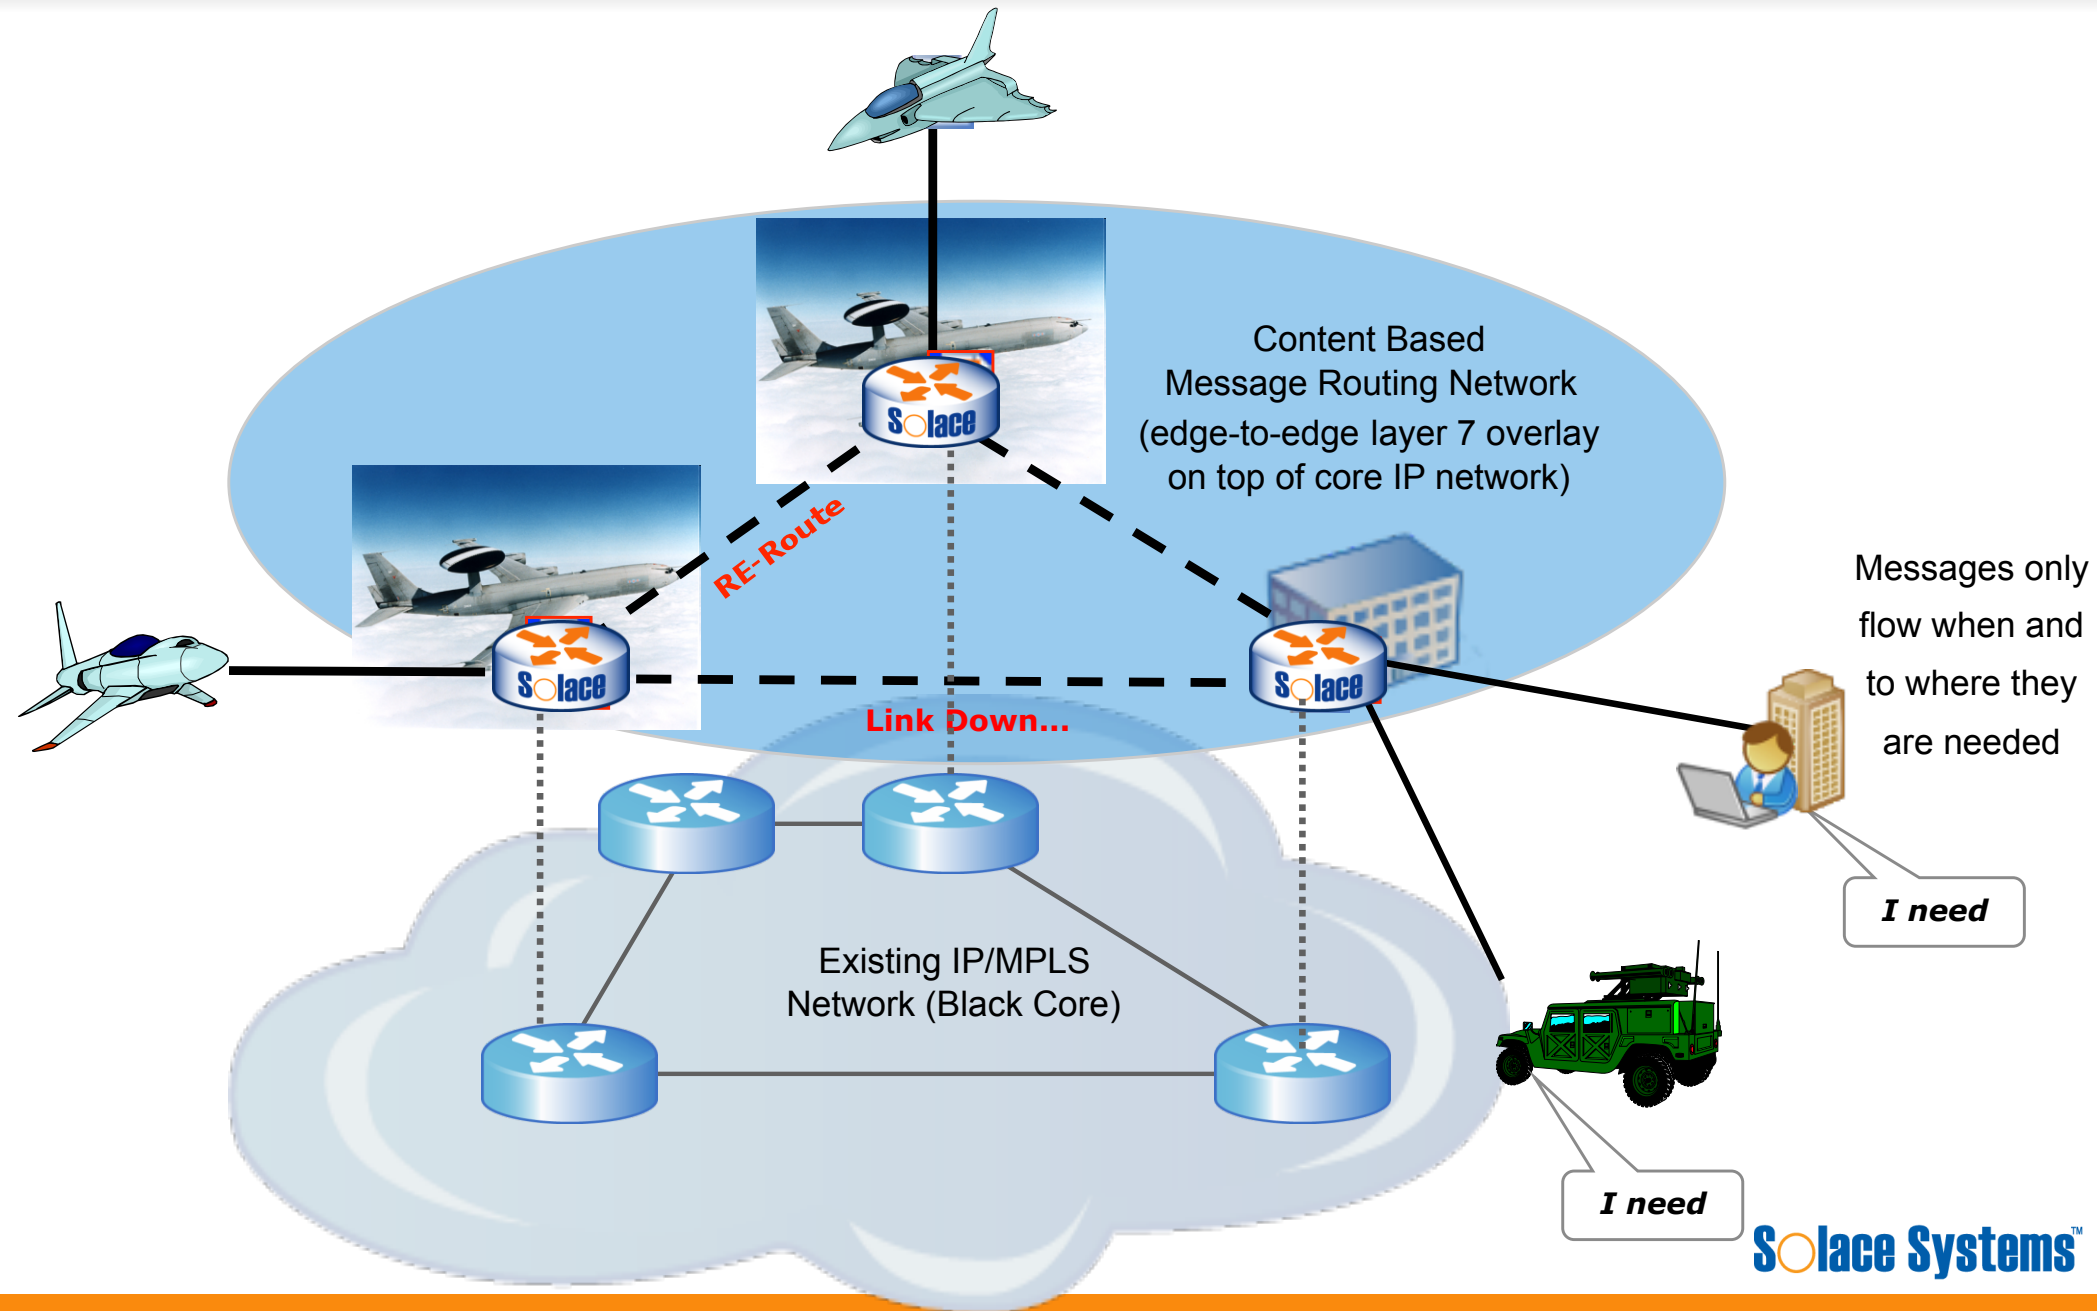


Geospatial overlay networks

- Geospatial content aware networking infrastructure
- Hardware-based messaging and middleware overlay for IP networks
- Automatic propagation of data by “point in polygon” subscriptions



“Sensor to Shooter” Message Distribution



BlockAway.com

craigslist

post to classifieds
my account
help, faq, abuse, legal

search craigslist

for sale

event calendar

S	M	T	W	T	F	S
24	25	26	27	28	29	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22

avoid scams & fraud
personal safety tips
craigslist factsheet
best-of-craigslist
job boards compared
lawsuit dismissed
craigslist movie & dvd
craigslist T-shirts
craigslist foundation
defend net neutrality

system status

terms of use privacy

boston gbs nwb bmw nos sob

community	housing	jobs
artists classes events general groups pets politics	activities childcare local news lost+found rideshare volunteers	accounting+finance admin / office arch / engineering art / media / design biotech / science business / mgmt customer service education food / bev / hosp general labor government human resources internet engineers legal / paralegal manufacturing marketing / pr / ad medical / health nonprofit sector real estate retail / wholesale sales / biz dev security skilled trade / craft software / qa / dba systems / network technical support transport tv / film / video web / info design writing / editing [ETC] [part time]
personals	for sale	services
strictly platonic women seek women women seeking men men seeking women men seeking men misc romance casual encounters missed connections rants and raves	barter bikes boats books business computer free general jewelry material rvs sporting tickets tools wanted	beauty computer creative erotic event legal lessons automotive
discussions	forums	gigs
1099 etiquet feedbk film fitness fixit	gifts haiku health help history housing jobs jokes kink l.t.r. legal linux loc pol m4m money motocy music npo open tax testing transg travel vegan w4w wed	crew event labor computer
	arts+crafts auto parts baby+kids cars+trucks cds/dvd/vhs clothes+acc collectibles electronics farm+garden furniture games+toys garage sale household motorcycles photo+video	creative domestic talent writing

us cities
atlanta
austin
boston
chicago
dallas
denver
houston
las vegas
los angeles
miami
minneapolis
new york
orange co
philadelphia
phoenix
portland
raleigh
sacramento
san diego
seattle
st bayarea
wash dc
more ...

us states
alabama
alaska
arizona
arkansas
california
colorado
connecticut
dc
delaware
florida
georgia
guam
hawaii
idaho
illinois
indiana
iowa
kansas
kentucky
louisiana
maine
maryland
mass
michigan
minnesota
mississippi
missouri
montana
n carolina
n hampshire
nebraska
nevada
new jersey
new mexico
new york
north dakota
ohio
oklahoma
oregon
pennsylvania
puerto rico
rhode island
s carolina

canada
calgary
edmonton
halifax
montreal
ottawa
toronto
vancouver
victoria
winnipeg
more ...

intl cities
amsterdam
bangalore
barcelona
berlin
buenosaires

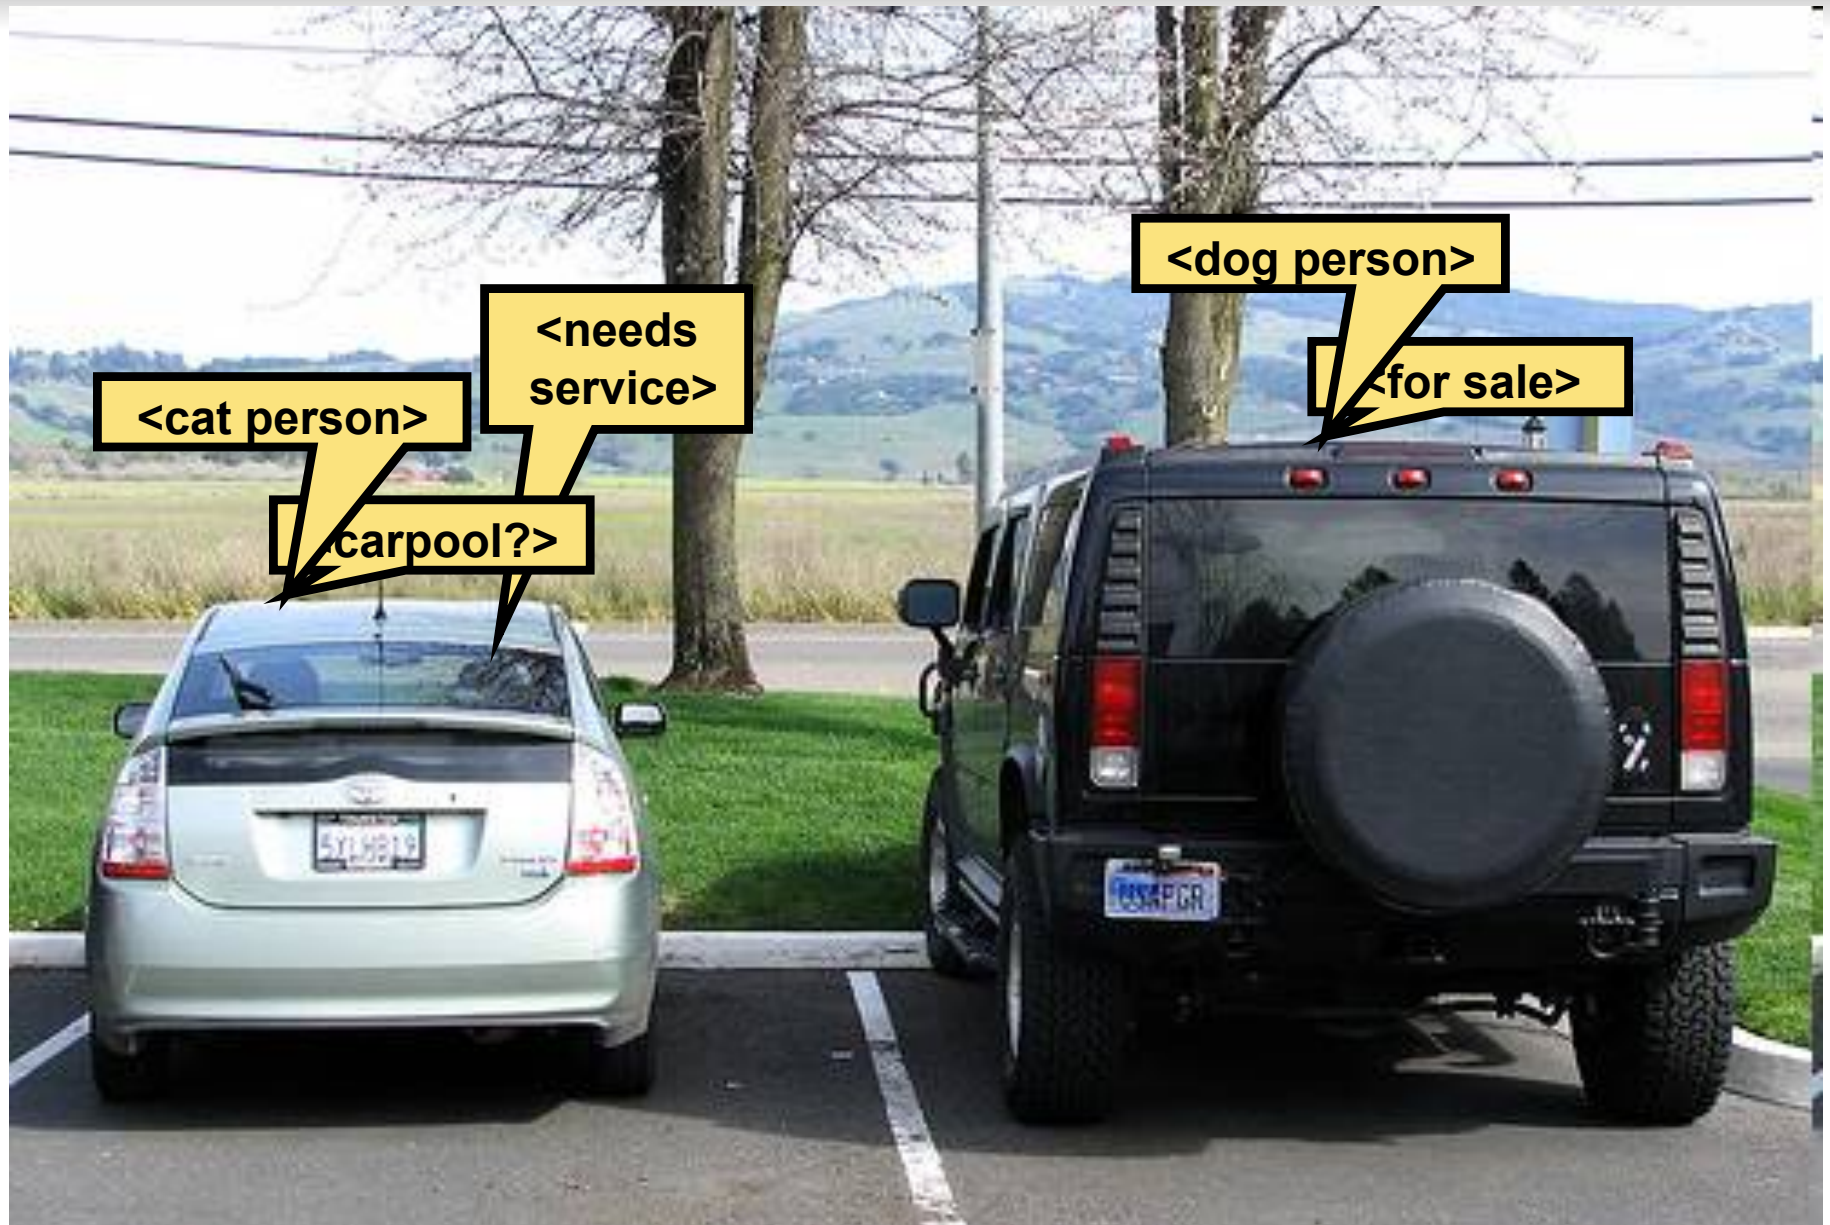
countries
argentina
australia
austria
bangladesh
belgium
brazil
canada
caribbean
chile
china
colombia
costa rica
czech repub
denmark
egypt
finland
france
germany
great britain
greece
hungary
india
indonesia
ireland
israel
italy
japan
korea
lebanon
malaysia
mexico
micronesia
netherlands
new zealand
norway
pakistan
panama
peru
philippines
poland
portugal
russia
singapore



Augmented Reality



Social Automotive Telematics



Online Gaming Use Case



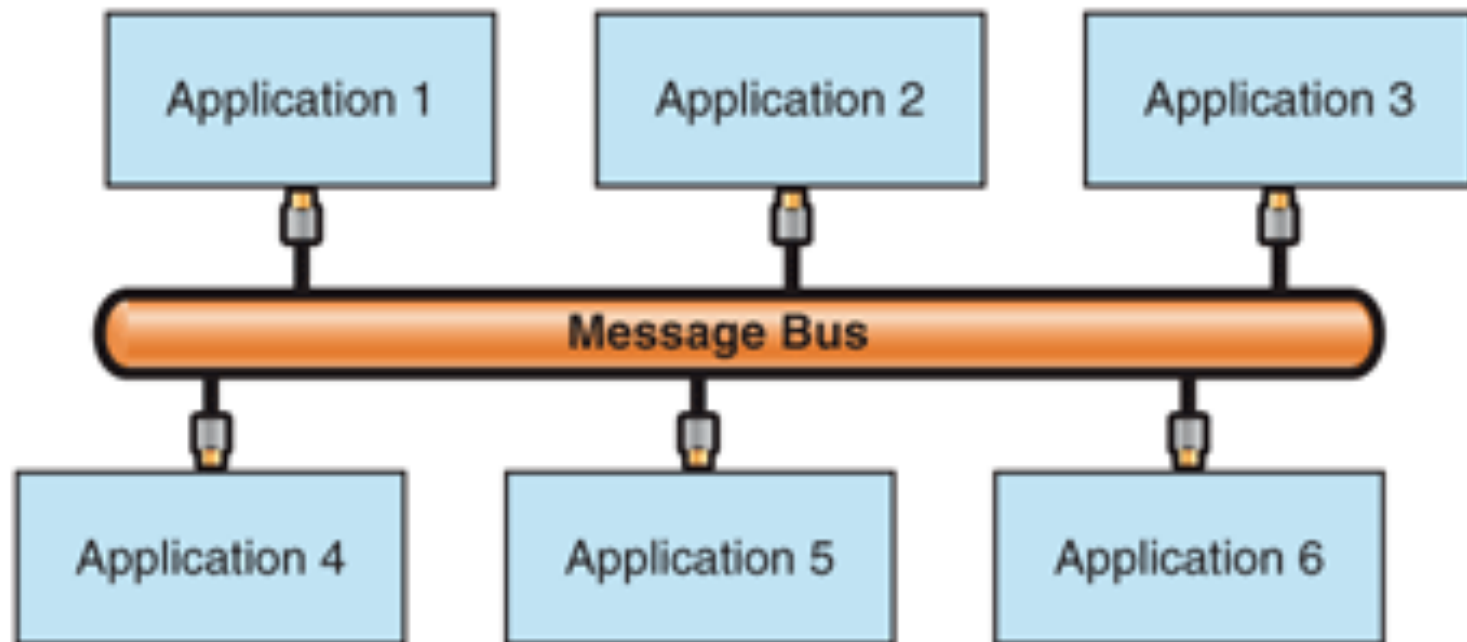
Existing Architecture



- UI events (cards & clicks)
- Bets
- In-game rewards
- Presence heartbeats
- Chat
- Financial Transactions

- Marketing events
- Tournament promotions
- Login/out events
- Statistical data
- Regulatory audit logging

Classic Job for a Message Bus?



Online Casino Company Requirements

○ Scalable

- Needed 150,000 msgs/s with guaranteed once in order delivery
- 1 million concurrent users growing to 10 million

○ Low Latency

- Web responsiveness with support for streaming to various RIAs

○ Reliability

- Both HA and DR failover required with no data loss


○ Operations

- Simplicity of appliance
- Flexible Management & monitoring

Superbowl Sunday Problem



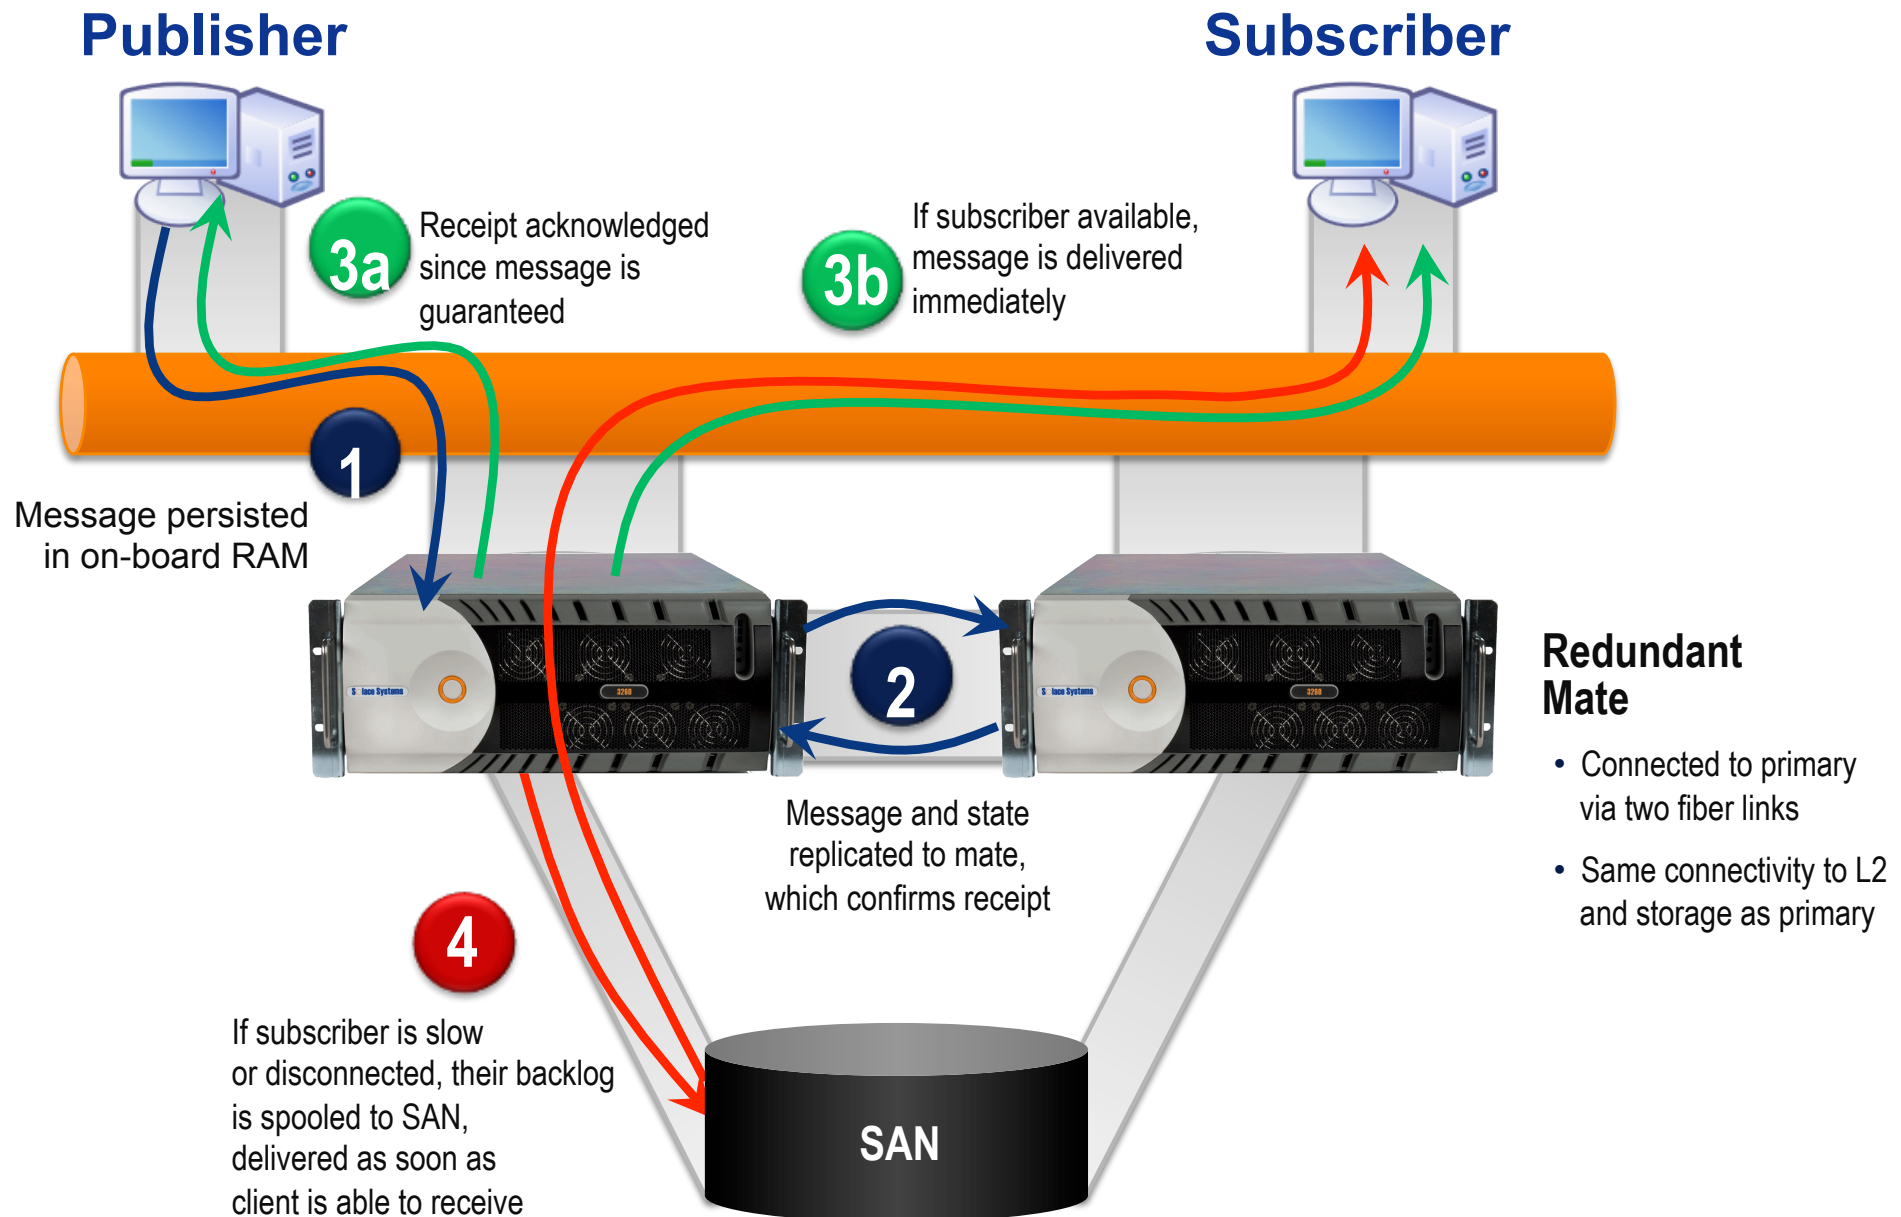
The Buzzer Problem



The Regulatory Problem


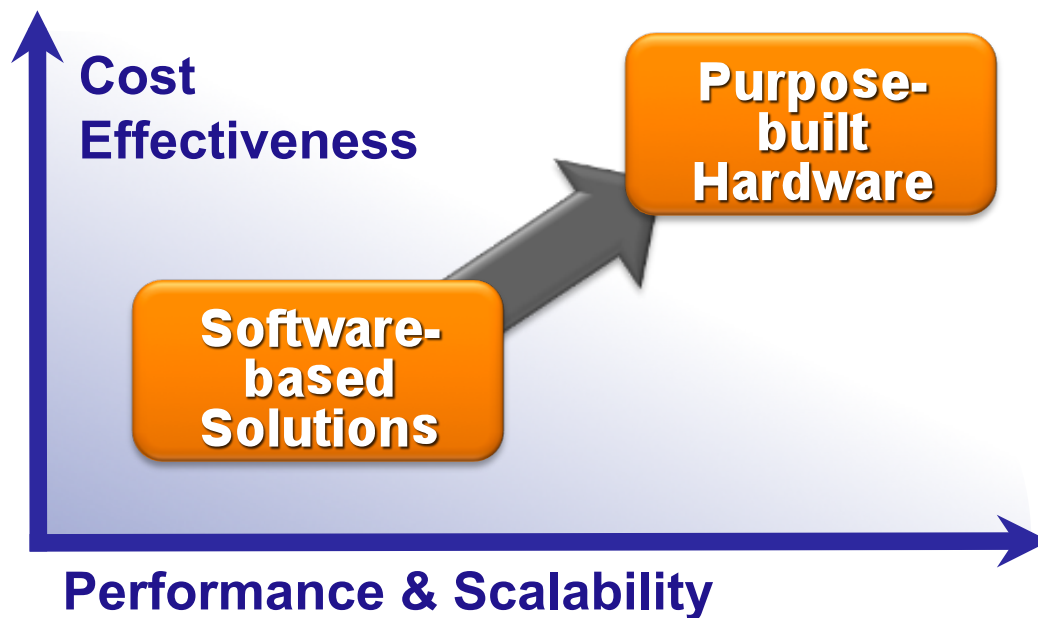


Fault Tolerant Clustering of Solace Routers



Inevitable Migration to Hardware

Well Understood, Critical Processing Paradigms
Always Migrate to Dedicated Hardware



What are you thinking of building?

