



Scott Healy
VP Systems Planning and Engineering
Sabre Holdings

Heterogeneous Air Shopping Computing Environment

14 April 2004

Why Heterogeneous Computing Environments?

The Drawbacks

- **More complex**
- **Operational headaches**
- **Troubleshooting challenges**
- **Additional staffing/skillsets to deal with complexity**



Why Heterogeneous Computing Environments?

The Benefits

Matching the platform to the workload produces better results

Objectives

Reliability
Scalability
Operability
Speed
Price/Performance
OLTP
Batch Processing

Choices

HP NonStop
Unix
Linux
NT
Open Source
3rd Party Solutions
Existing Apps

Results

Lower cost
Better performance
Adequate reliability
Migration paths
Hardware
Programming Languages



Sabre Holding's evolution to a heterogeneous air shopping environment

- **About Sabre Holdings**
- **Business drivers and rationale for Air Travel Shopping Engine (ATSE) Hybrid**
- **HP NonStop's role in the computing environment**
- **HP Linux role in the computing environment**
- **System architecture**
- **Technical challenges**
- **Hybrid in production**
- **Where do we go from here?**



Sabre Holdings' Business Units

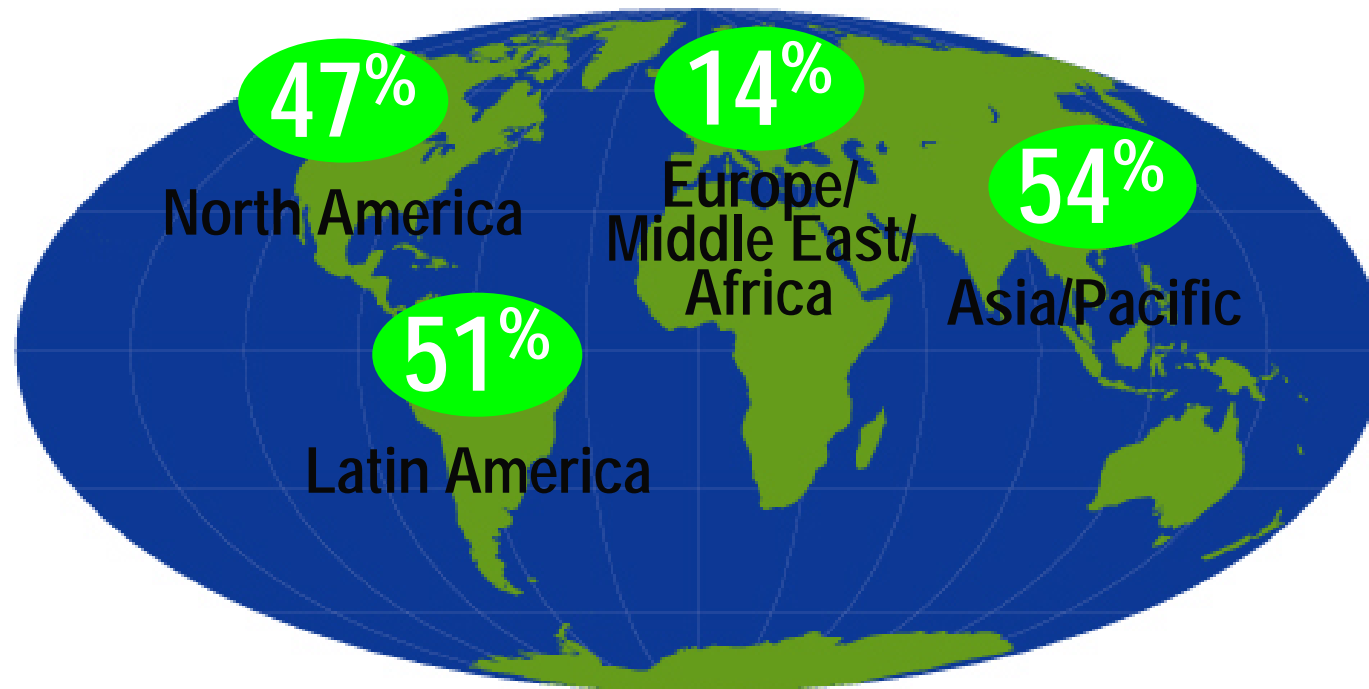


Sabre Holdings' Business



- **Leading provider of travel commerce technology**
 - Headquarters in Dallas / Ft. Worth, TX
 - 6,500 employees in 45 countries
 - 2003 revenues of over \$2 billion
 - S&P 500 corporation

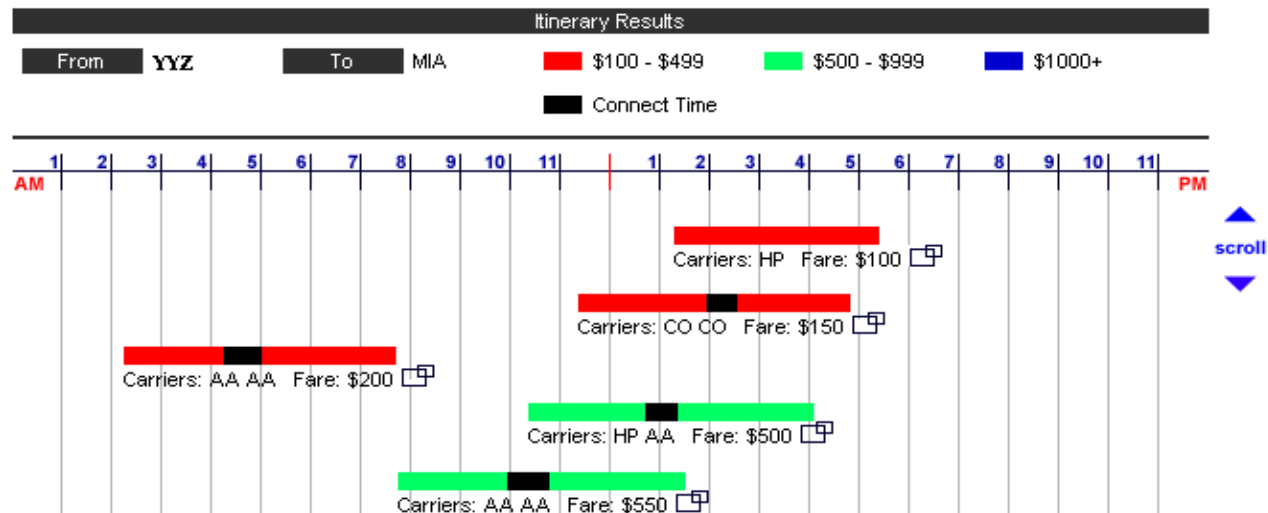
Worldwide Distribution Market Share



Air Shopping Business Problem

Search multiple airlines, flights, fare types, dates and adjacent airports in order to find desirable itineraries.

Calculate exact price, taxes and surcharges for ticketing, refund and exchange purposes.



Healy's law of air shopping:

You can't make it cheap enough or fast enough



Air Travel Shopping Engine

- **Objective – Multi-year plan to move Air Shopping and Pricing off IBM mainframes to open systems environment**
- **Phase I – Domestic US**
 - Start on NSK, investigate heterogeneous environment
 - Add Canada in Phase 1.5
- **Phase II – International / new shopping models**



Business Drivers: Cost, Reliability, and Performance

- **Air Shopping demand for large numbers of processors**
 - High CPU consumption for the target applications
 - Average 3+ CPU seconds/transaction on S86000
 - Large variance among transactions
 - Ever increasing volume and complexity
 - Large memory footprint for the application
 - High wait times for external data
- **No direct revenue for transactions**
- **Competition requires staying on the “leading edge”**
 - Database must be highly reliable and scalable
 - Transaction processing must be scalable and fast
 - Must keep up with technology - Twice as fast for half the cost every 18 months....



Heterogeneous Approach

Deploy commodity Linux servers alongside HP NonStop servers to create a more cost-effective platform for each ATSE workload

- HP NonStop delivers high availability and reliability
 - Better than or equal to TPF at significantly lower cost
 - Best fit for critical workloads and master database management
- Linux delivers 64-bit memory model and faster CPUs
 - Lower availability and reliability than HP NonStop but at significantly lower cost
 - Best fit for CPU-intensive shopping workloads



HP Nonstop Applications

- **Master database server**
 - Schedules
 - Fares and rules
 - Availability Status (AVS)
 - Direct Connect Availability (DCA) cache
- **Critical application server**
 - Guaranteed air itinerary pricing
 - Enhanced availability – (schedules and seats)
 - Fare quotation, rules display*

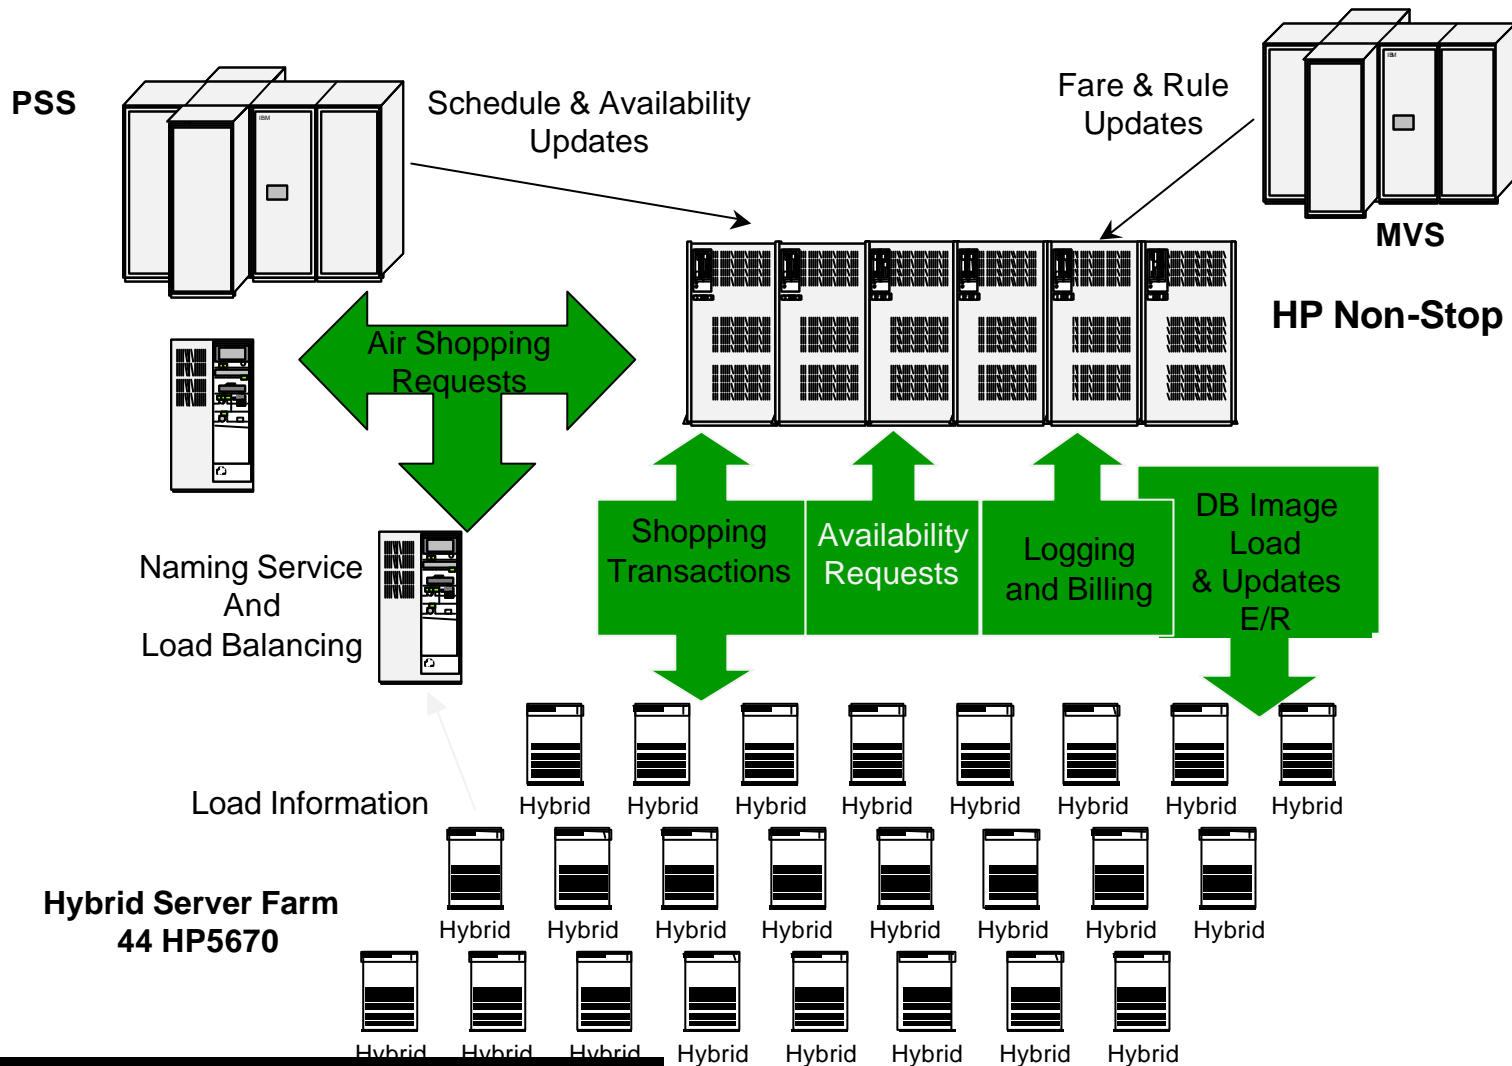


HP/Linux Role

- **Horizontally scalable, low cost computing platform**
- **Database copies from NSK**
 - Read only high volume data replicated to HP/Linux
 - Highly volatile inventory data stays on the NSK
- **Air shopping transactions are stateless**
- **Nothing on the HP/Linux system needs to be recovered if the box fails**
 - Low cost and performance drives software selection
 - Planned over capacity (n+2) provides fault tolerance
 - Billing & logging data written back up to NSK



System architecture



Hybrid System Software

- **Extensive use of open source software**
 - MySQL 4.0.15
 - GNU compilers – g++ 3.2.3 and glibc 2.3.2
 - TAO object request broker
 - Redhat Linux AS 3.0
- **3rd party software**
 - Golden Gate Extractor/Replicator
 - Monitoring – Prognosis
 - Security
- **Sabre developed applications and scripts**



Technical Challenges

- **Porting code from HP NonStop to UNIX & Linux**
 - 1 tech, 1 day (not a challenge)
- **Clearcase source code management for 64 bit Linux**
 - Opted to use a cross-compiler versus exported views
- **Load balancing and distribution**
 - Meridian load balancing distributes load as workload increases
 - Investigating load balancing alternatives

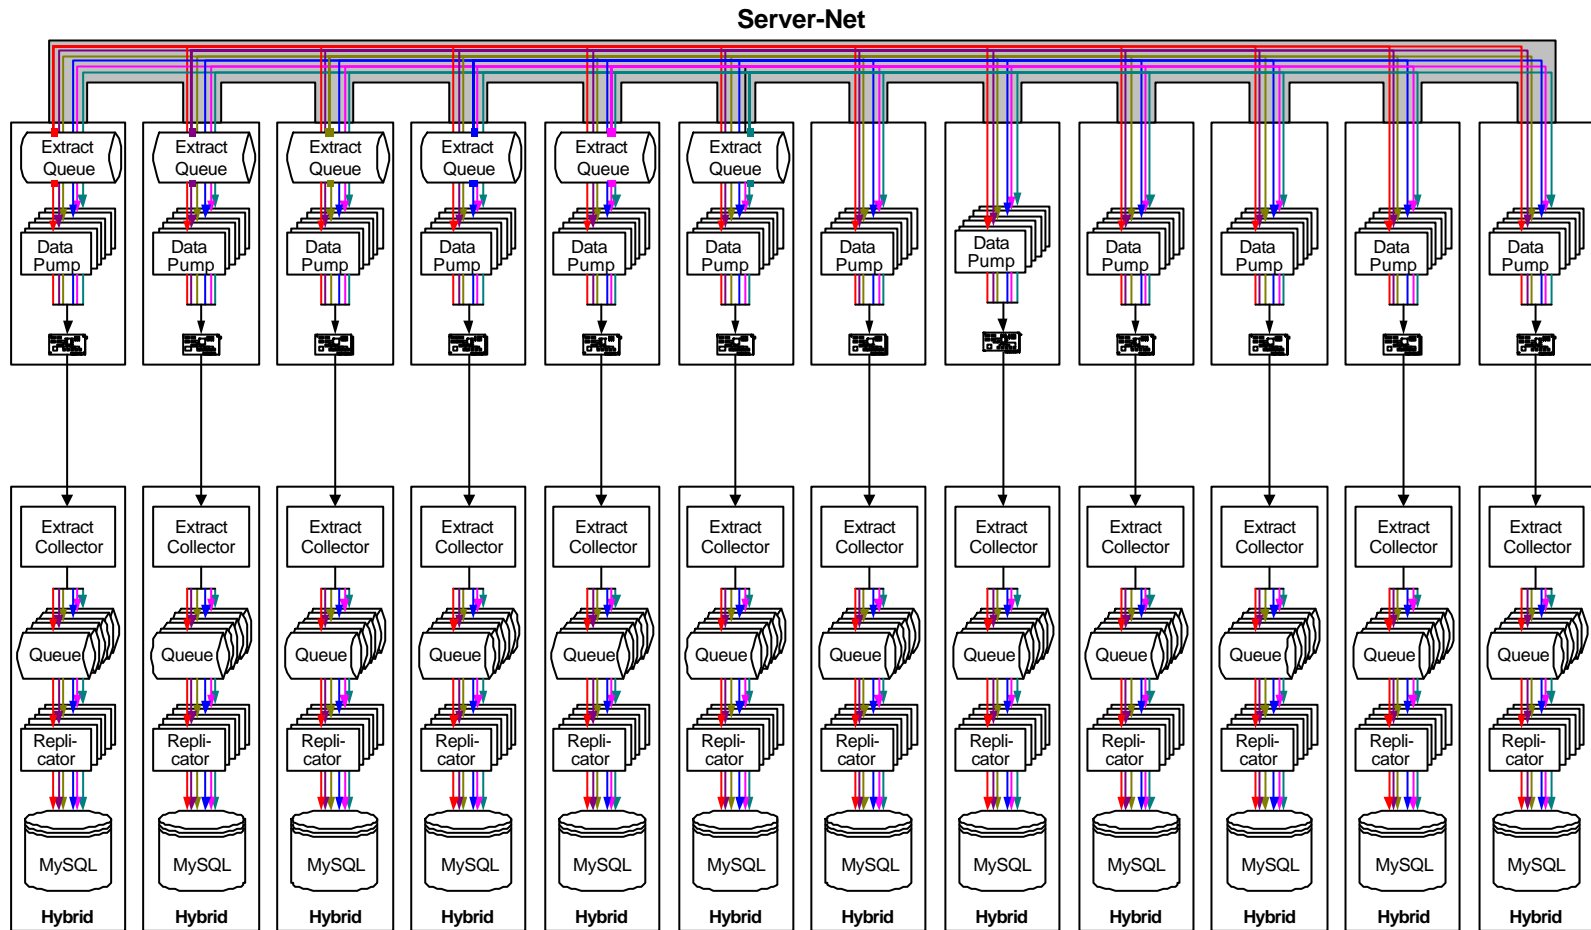


Technical Challenges - Hybrid Database

- **Data transport to the hybrid farm**
 - Extensive use of application cache for performance
 - Bandwidth and I/O overhead was a concern
 - DBMS added to hybrids to reduce transport load
- **Data replication to the hybrid DBMS**
 - Bandwidth was an issue for scalability
 - Normal replicate rate is 300,000 updates/hour
 - Used discrete event simulation model to determine max hybrid limit
- **Through consultation with Golden Gate, determined use of E/R data pumps would provide intermediate term scalability**



Data Replication



Technical Challenges - Hybrid Database

- **The hybrid MySQL database maintains a live database with updates replicated from the NonStop master.**
 - The database must be active and available on all Hybrid servers accepting traffic.
 - Downtime for the database falls under the same criteria as that for hardware failures.
- **The hybrid MySql database contains approx 150 tables, NonStop has approx 280.**
 - Updates are truly 24X7
 - Characterized by spikes of updates corresponding to batch transmissions from data vendors and airlines



Hybrid in Production

- **Sabre has been running air shopping transactions on both NonStop and HPUX since July 2003**
 - First configuration used HP 5405s
 - Performance of hybrids for air shopping transactions was similar to NonStop 86000s
- **Sabre upgraded to Linux HP 5670s in October 2003**
 - 95% of low fare search traffic on hybrid farm
 - Better response time
 - NSK serving cached carrier availability data for Shopping requests



Where do we go from here?

- **Improve reliability of overall system**
- **Increase system capacity through “application only” hybrid servers linked to MySQL servers**
- **Migrate additional workloads to the hybrid farm**
- **Expand the DCA cache on NSK**
 - Proactive polling
 - Dynamically determined expiration times
 - Charging from additional sources
- **Continue to evaluate future requirements**



Summary

- **Business cases can readily be built for creating heterogeneous environments**
- **Challenges exist but can be overcome**
- **Companies that master the heterogeneous environment will be rewarded with better performance and lower costs**
- **Lots of companies are doing it...**



Questions?

