

HP Enterprise Data Warehouse Deep Dive

Steve Tramack, Sr. Engineering Manager, I2A Solutions, HP

Microsoft



HP and Microsoft Strategic Partnership

Igniting the contribution of IT to the business



Microsoft | **HP**

World-leading software and hardware companies

- 3-year, \$250M joint investment
- Portfolio of solutions delivers rapid time to value
- Integrated HP and Microsoft Services
- Bridge the path to the cloud




Microsoft

Converged Systems deliver bigger business impact

Deploy in days or
weeks, not months

Boost
application
performance
on balanced
systems



Deliver with
confidence and
ensure project
success

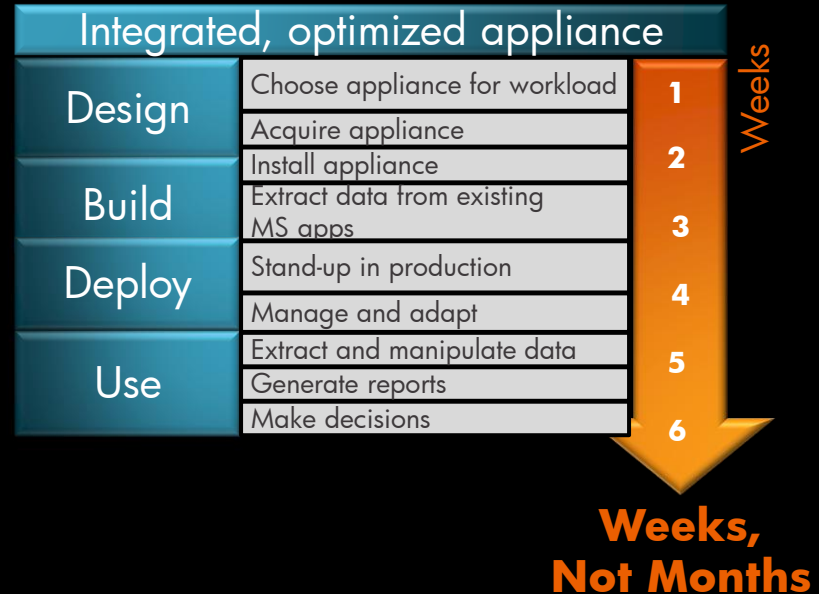
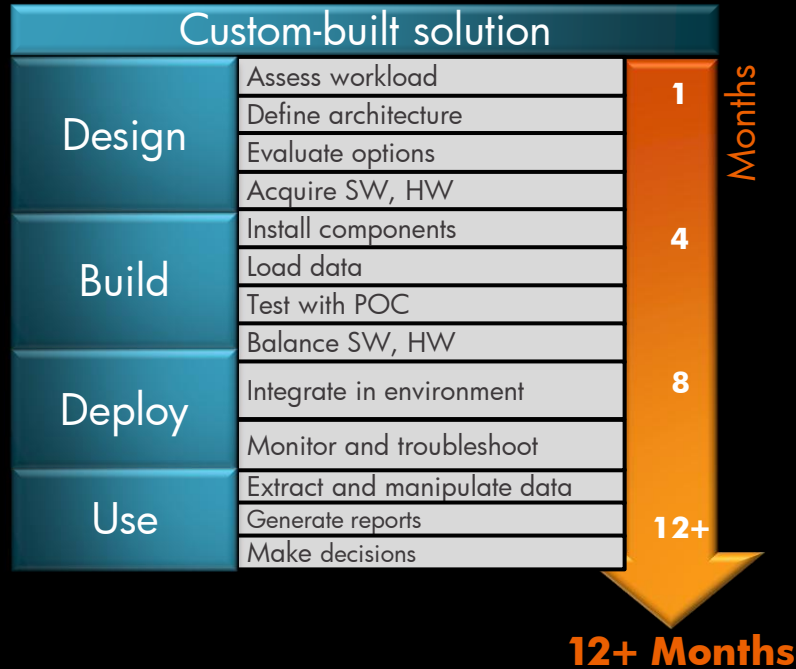
Streamline management
and lower TCO



Microsoft

Accelerate time to value

Deploy in weeks rather than months



Microsoft

DATA WAREHOUSING LANDSCAPE AND OFFERINGS



Microsoft

Data Warehouse Landscape and Customer Challenges

- 1 Data explosion – number of firms with 10+TB expect data will double in 3 years
- 2 Customers need to reduce cost
- 3 More customers plan to do real-time analytics
- 4 Growing market adoption of DW appliances
- 5 Move to MPP – more customers plan to use MPP in the next 3 years

1



17%

34%

2



57%

3



92%



4


Appliances

78%

5

Massive Parallel Processing

82 %



Growing Market

HP and Microsoft Data Warehousing Offerings

Tier 1 Offerings


HP Business Data Warehouse Appliance	Microsoft® SQL Server® 2008 R2	Microsoft® SQL Server® 2008 R2 Fast Track Data Warehouse	Microsoft® SQL Server® 2008 R2 Parallel Data Warehouse
An affordable SMP solution for data warehousing on optimized hardware	Scalable and reliable platform for data warehousing on any hardware	Reference architectures offering best price performance for data warehousing	Appliance for high-end data warehousing requiring highest scalability, performance, or complexity
Ideal for small data marts up to 5 TB	Ideal for data marts or small to mid-sized enterprise data warehouses (EDWs)	Ideal for data marts or small to large data warehouses with scan-centric workloads	Offers flexibility in hardware and architecture
Appliance	Software only	Reference Architecture (software and hardware)	Appliance
Scale Up Data Warehousing	Scale up data warehousing	Scale up data warehousing	Scale out data warehousing with massively parallel processing (MPP)
Up to 5TB	10s of terabytes	4–80 terabytes	10s–100s of terabytes
\$103K (includes hardware)	\$27.5K per processor \$8.6K per server \$164 per CAL \$57.5K per Proc only (Datacenter)	\$107K–\$683K (2–4 Procs; includes hardware)	\$1.5–\$1.7 million per rack (includes hardware) \$13K – 44K per terabyte

Tier 1 Services and Support



The Appliance Engineering Approach

Each solution matching and balancing four main elements

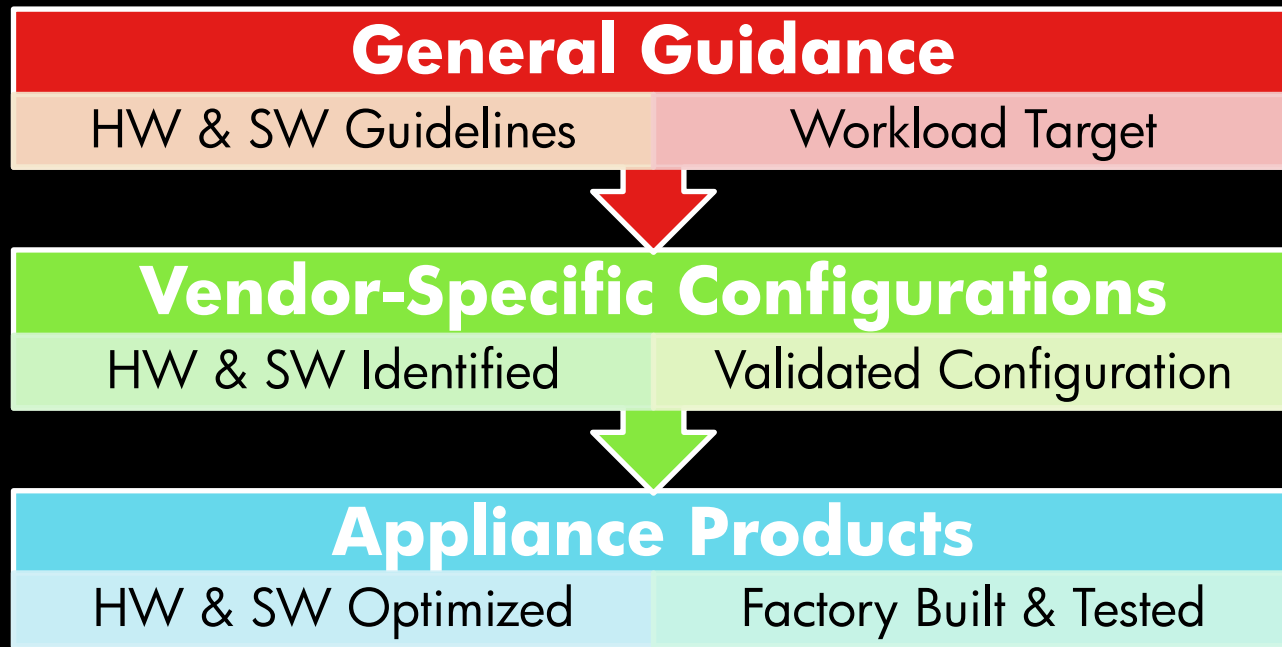


- Driven by a fundamental understanding of the workload
- Architecture followed by and supported by components



Microsoft

Reference Architecture Maturity Model



Successfully Design
a new system using
best practices

Successfully Buy
a new configuration
with known abilities

Successfully Deploy
a new solution with
proven performance



Microsoft

HP & Microsoft Data Warehousing Continuum:

SQL 2008 R2 Server Fast Track, HP Enterprise Data Warehouse Appliance

Scales from 1TB to Hundreds of TBs

- ❖ Balanced solutions ideal for data marts - EDW with scan-centric workloads
- ❖ Packaged and custom support
- ❖ From SMB to Enterprise
- ❖ Built on HP ProLiant G7



Reference Architectures

Business Data Warehouse

ProLiant DL370 (2P)
Up to 12 Cores
Internal HDD (6TB)

Basic RA

ProLiant DL38x (2P)
Up to 24 Cores
P2000 G3
(up to 20TB)

Mainstream RAs


ProLiant DL58x G7 (4P)
Up to 48 Cores
P2000 G3
(Up to 60TB)

Premium RA

ProLiant DL980 G7 (8P) Up
to 64 Cores
P2000 G3
(Up to 80+ TB)

HP Enterprise Data Warehouse

Per /rack: 11xProLiant DL360 (2P)
10 x P2000G3 (56 – 150TB)
Up to 4 data racks (600TB)



Appliance

Tier 1 Enterprise Data Warehouse

Value to Customer

- Enterprise-class scalability to 100s of terabytes
- High performance
- **Interoperability** with leading BI products
- Tier 1 support and maintenance
- Mature SQL Server platform with **high security and robust engineering process**
- Strong data warehouse **vision and roadmap** that includes industry leading technologies



Supporting Features

- MPP with **shared nothing architecture**
- Distributed query optimization
- Balanced hardware with pre-tested and **pre-tuned appliances** optimized for data warehousing
- third-party product **integration** (for example, Microstrategy, Business Objects, Informatica)
- **Tier 1 support** and maintenance
- Roadmap includes column store, petabyte scalability, real-time data warehousing, MDM, and Data Quality



Microsoft

Affordability & Choice

Value to Customer

- Low Cost:
 - Low hardware cost
 - Low maintenance cost
 - Less risk and disruption on hardware standards
- Choice of deployment options
 - Through distributed architecture that supports MPP hub and SMP spokes
 - Enables customer to support users with different needs and SLAs



Supporting Features

- Broad range of options with 1 rack to 4 rack, plus Dev/Test appliances
- Use of industry standard hardware
- Lower acquisition and maintenance cost than Teradata, Oracle, and IBM
- Distributed architecture supports PDW as the Hub and Fast Track data marts as spokes



Microsoft

The Most Complete Data Warehouse Solution



Value to Customer

- Compelling End-user BI experience
- Agility to support new business groups with integrated data marts
- Fast time to value for BI through integration with Microsoft BI
- Complete data warehouse solution spanning desktop, data marts, and EDW

Supporting Features

- PowerPivot, Reporting Services, and Microsoft Excel®
- Distributed Data Warehouse Architecture integrates both MPP and SMP data warehouses
- Integration with Microsoft BI tools, including Analysis Services, Integration Services, and Microsoft SharePoint®
- Most comprehensive toolset, including ETL, BI, MDM, and StreamInsight



Microsoft

HP ENTERPRISE DATA WAREHOUSE APPLIANCE




Microsoft

HP Enterprise Data Warehouse Appliance


Transforming SQL deployment

BEFORE



- No single view of data across the enterprise
- Multiple siloed SQL data marts
- Proprietary, expensive EDW appliances

AFTER



- The world's most scalable, easy-to-manage enterprise data warehousing solution
- 600TB+ = 60x Library of Congress printed collection




Microsoft

HP Enterprise Data Warehouse Appliance

Optimized for SQL Server 2008 R2 Parallel Data Warehouse

Cut the complexity and expense of deploying very large data marts with industry-leading price/ performance and mission-critical support



Complete

- All-in-one data warehouse appliance for mission-critical environments; New consulting services for HP Enterprise Data Warehouse



Optimized

- Single view of information across the enterprise
- One-step solution update
- Enterprise-class scalability for complex queries at market-leading \$/TB



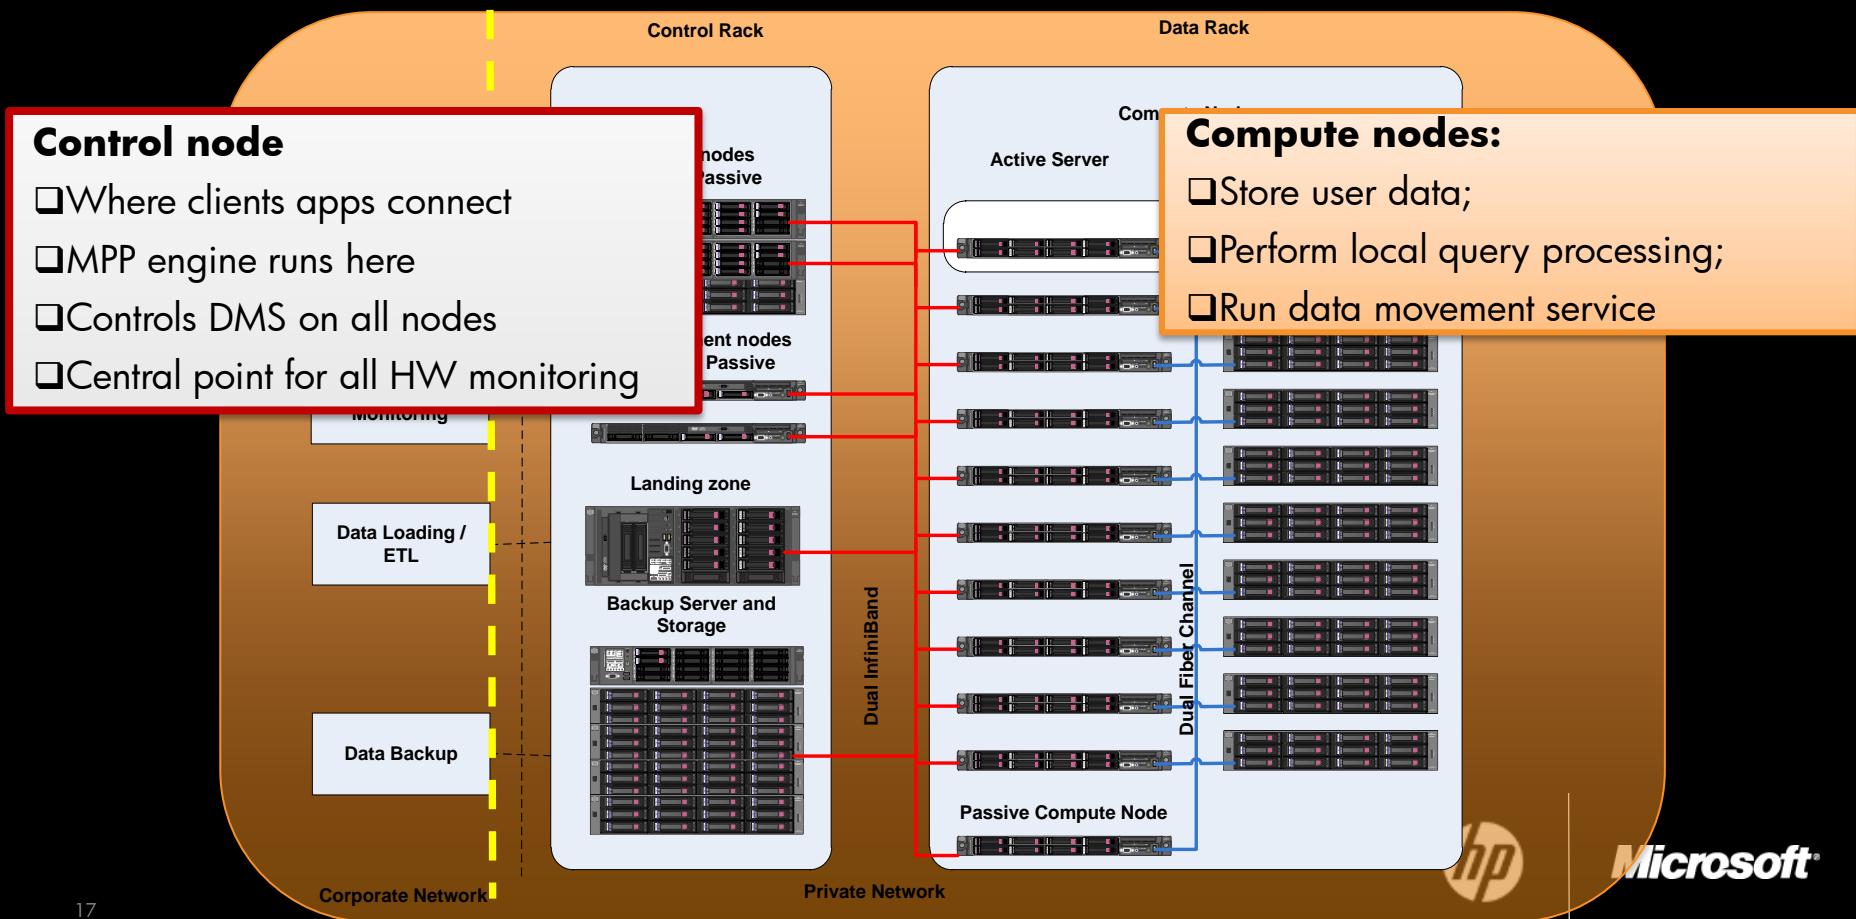
Agile

The world's most scalable, easy-to-manage enterprise data warehousing solution
Offers 8x scale to over 600 TB, and 40x to 200x faster queries



Microsoft

Enterprise Data Warehouse Architecture



Enterprise Data Warehouse Components

Data Racks
10 + 1 compute
nodes per rack



Enterprise Data Warehouse HA Features

Failover
Clusters

Redundant
fabrics

Redundant
PDUs



Redundant
power

Redundant
cooling

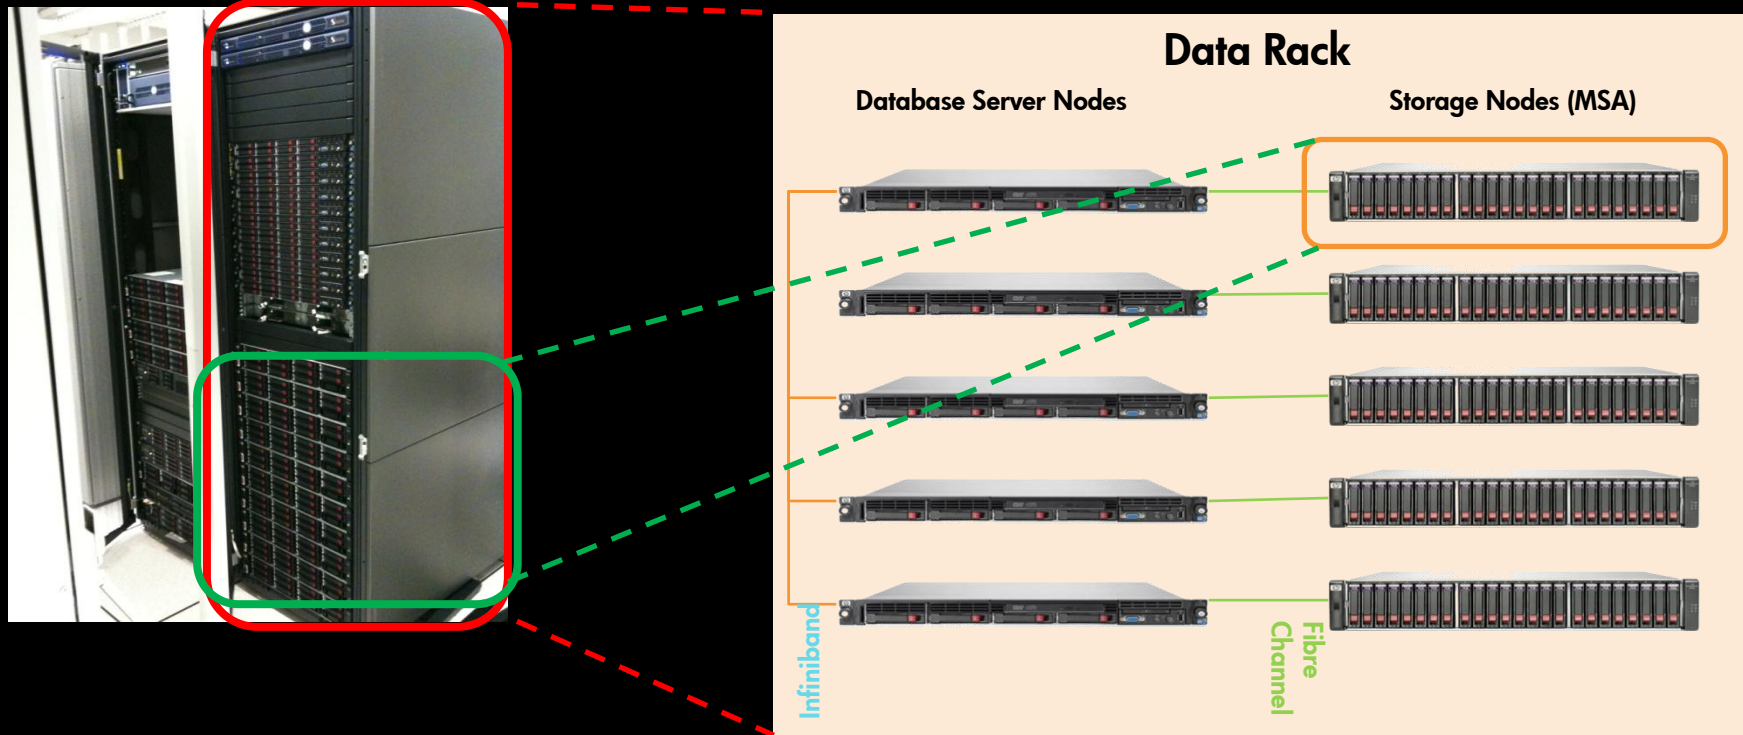
RAID 10 and
hot plug disks

Pre-failure
alerting



Microsoft

High-Level PDW Architecture



Microsoft

Architecture

Storage Node Anatomy



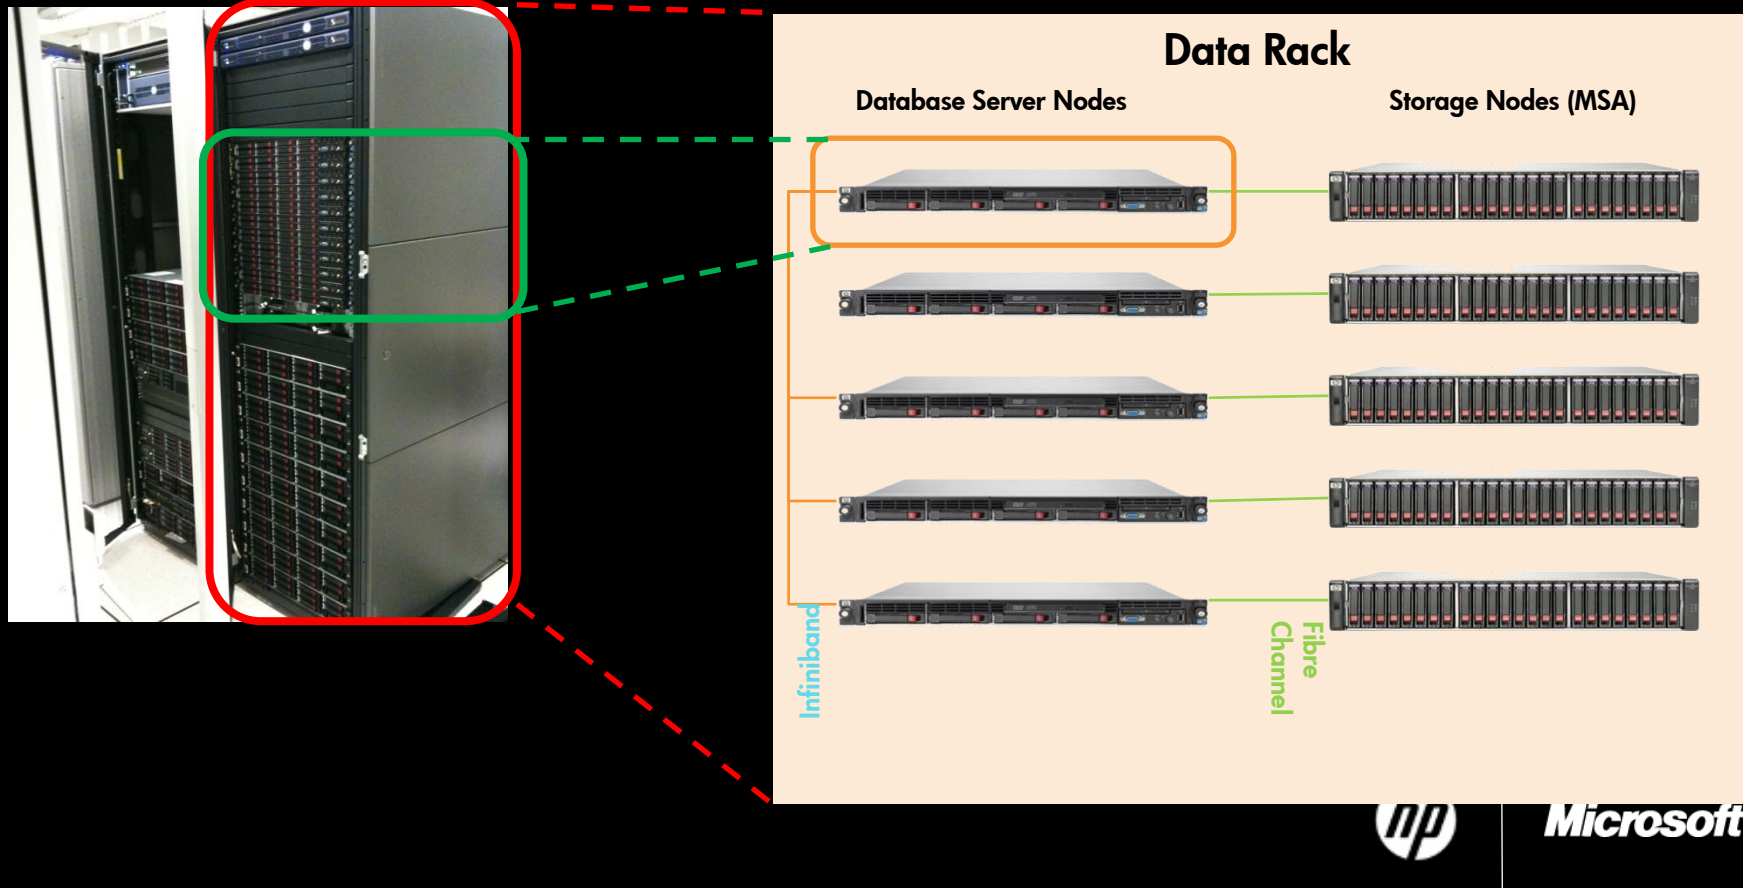
MSA P2000 G3

- 11x HP P2000 **300GB** 6G SAS **15K** 3.5in ENT HDD, OR
- 11x HP P2000 **1TB** 6G SAS 7.2K 3.5in MDL HDD, OR
- 24x HP **300GB** 6G SAS 10K 2.5in DP ENT HDD, OR
- 24xHP **600GB** 6G SAS 10K 2.5in DR ENT HDD



Microsoft®

High-Level PDW Architecture



Database Server node



- DL360 G7 (X5670)
 - 96GB RAM (12x HP 8GB 2Rx4 PC3-10600R-9)
 - HP DL360 SFF HD Backplane
 - HP 512MB P-Series BBWC Upgrade
 - 8x HP 300GB 6G SAS 10K 2.5in DP ENT HDD
 - HP IB 4X DDR PCI-e DUAL PORT 0 mem HCA
 - 2x HP 460W HE 12V Hotplug Power Supplies
 - HP iLO Advance license
 - 6 internal tempdb disks
 - 2 OS disks



Microsoft

Measured Performance and Scalability

Metric	Measurement
IO Scan Rate	1.65GB/sec per MPP node; 16.5GB/sec per data rack <i>(12 x 1TB LFF MDL SAS: 10.7GB/sec/rack. Load and backup rates similar to SFF)</i>
SQL Scan Rate	1.36GB/sec per MPP node; 13.6GB/sec per data rack
Capacity per data rack	76TB (300GB SFF) to 152TB (600GB SFF)
Load rates	2.3TB/hour
Backup rates	5.8TB/hour to 12TB/hour (depending on compression rates)
Restore rates	3.8TB/hour to 10TB/hour (depending on compression rates)

System Under Test: 1 Data Rack EDW with SFF SAS disks. AU2 version of PDW
Workload: Industry standard DW workload

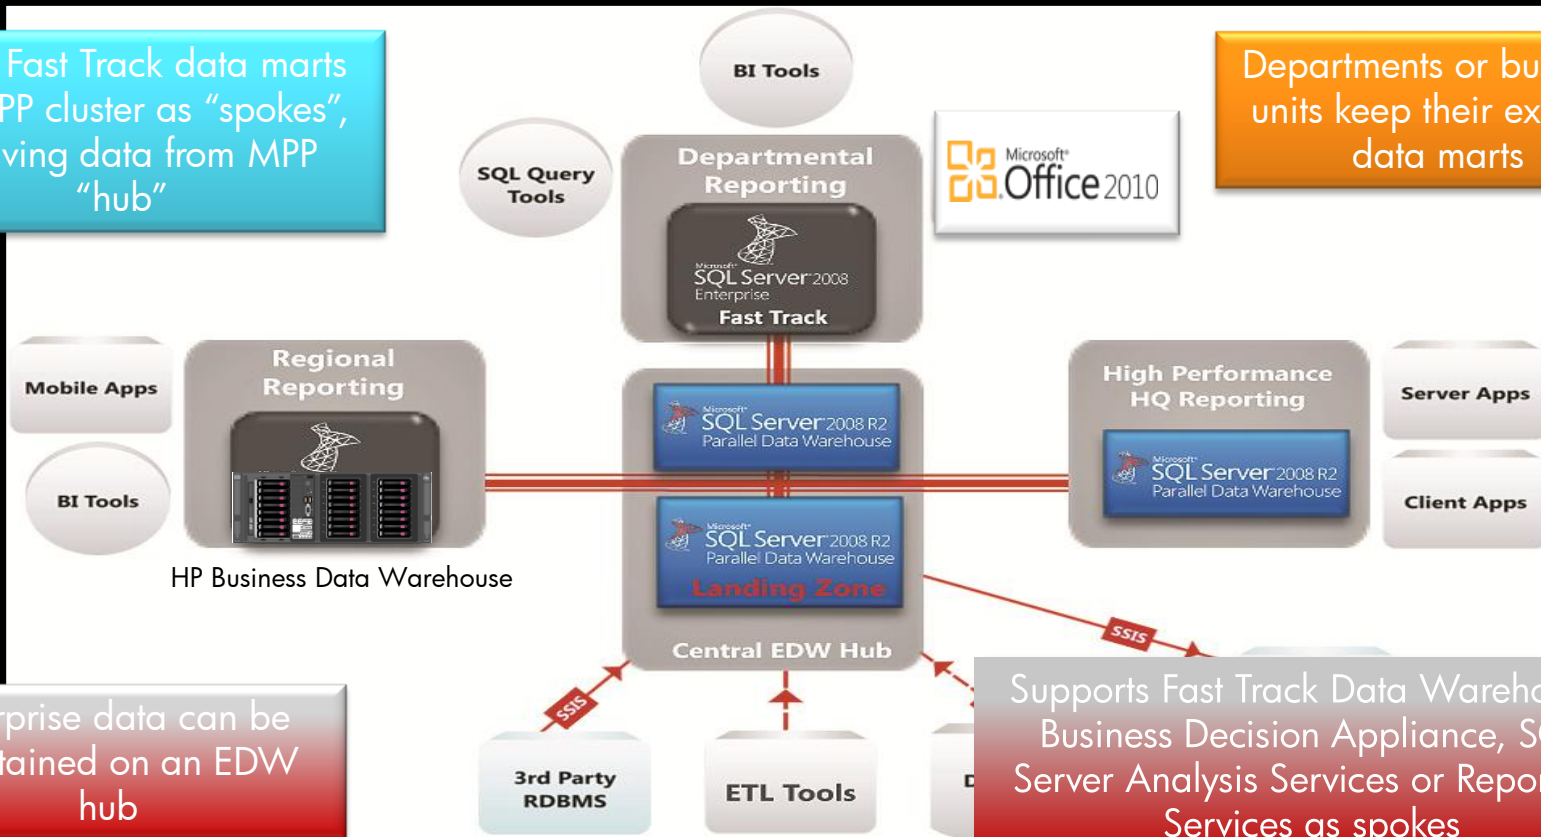


Microsoft

Parallel Data Warehouse Hub and Spoke

"Plug" Fast Track data marts into MPP cluster as "spokes", receiving data from MPP "hub"

Departments or business units keep their existing data marts



Enterprise data can be maintained on an EDW hub

Supports Fast Track Data Warehouse, Business Decision Appliance, SQL Server Analysis Services or Reporting Services as spokes

HP Enterprise Data Warehouse

Case Study

Challenge

- Need better storage and access for report generation, balance checking and data analysis

Solution

- Chose HP Enterprise Data Warehouse Optimized for SQL Server 2008 R2 Parallel Data Warehouse
- Quick to install

Results

- Reduced implementation risk
- Delivered 3X the performance at 1/3 the cost
- Can now generate reports in 15 minutes instead of four hours



"Our expectations for setup, ease of management, and processing speed were exceeded at every point by HP and Microsoft...we have tested the solution to generate reports in under 15 minutes for what used to take over four hours.".

Rich Hochron, CTO, Direct Edge



Microsoft

"Our success to manage the business for more than 220 retail stores depends on our ability to gain insight into those stores to make the most optimal decisions, so that we can proactively manage fluctuations in demand, enhance operations and bolster business growth. Microsoft Parallel Data Warehouse based HP Enterprise Data Warehouse (EDW) Appliance's turn-key approach, coupled with HP and MSFT's best practices allowed us to deploy the EDW Appliance in a few days. It only takes 3 hours to load data, which used to take more than 7 days to load. We are able to make business decisions significantly faster as report generation and analysis of the information is performed in a more timely manner."

- Assistant VP of Data Warehousing, Large US-based retail chain



Microsoft

EDW Support Pack

- Installation tools
 - Default password update tool
 - Time/Date synch tool
 - BIOS setting synch tool
- Troubleshooting tools
 - EDW Hardware Diagnostics tool
- Ongoing operational tools
 - MRA Compliance Checker
 - Push Button Update Utility (for firmware and drivers)



Microsoft

DEMO: Push Button Update Capabilities



Microsoft

Dramatic Simplification

Game-changing solutions from HP and Microsoft



2,000 components updated
with **one** button



HP Enterprise Data Warehouse Appliance


Consistent, validated and up-to-date environment



Microsoft

Converged Systems are Built on Converged Infrastructure.

The right foundation for the workload



Joint HP and Microsoft services for...

Seamless support:

Collaborative partnering of HP
Services with Microsoft
Services

Consulting expertise:

Expertise with on-premises,
outsource, or cloud
delivery models

Comprehensive offering:

Over 23,000 combined HP and
Microsoft professionals worldwide
deliver a complete lifecycle



Microsoft

Best-of-breed Partnership

Delivering faster time to value

HP + Microsoft
co-innovation



Integration of Converged
Infrastructure through
application services



R&D + Services + Partner
investments



Microsoft



THE **INSTANT-ON**
ENTERPRISE IS HERE.