



Embracing the BI Appliance wave

November 2010

Experience the Power of Knowledge™

Corporate Headquarters | Saama Technologies, Inc. | 900 E Hamilton Avenue, Suite 120 | Campbell, CA 95008 USA



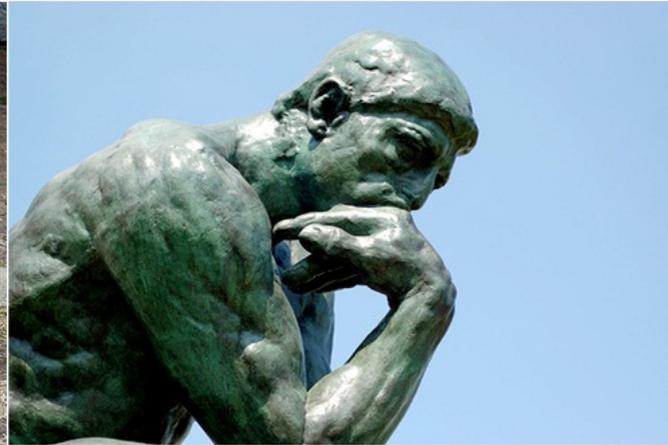
Agenda

- **About Saama Technologies**
- **What is a BI Appliance?**
- **Implementing a BI Appliance - *A Case Study***
 - *Shortlisting process*
 - *Selection process*
 - *Implementation process*

About Saama



Focused — Trusted — Experienced — Passionate



- Founded 1997
- Focused, pure-play BI provider
- Unique, innovative IP-based solutions
- Experienced leadership team
- 350+ experienced consultants
- Senior Board of Advisors
- Strong strategic partners
- Blue-chip client list
- HQ – Campbell, CA
- Offshore Development – Pune, India

About the Presenters



Ashish Mirji
Associate Principal



- 13 years in the IT services industry and specialization in Data Management and Business Intelligence.
- Worked as a Data Modeler, Datawarehouse Architect and successfully led several Data Management initiatives for several clients
- Engaged with clients like Genentech, General Electric, Franklin Templeton, Emirates Airlines, CIT, etc.
- Currently engaged as a Senior Manager to build Saama's Life Science practice

Winston Pinto
Engagement Manager



- 10 years in the IT services industry with specialization in Business Intelligence.
- Worked on BOBJ and Pentaho development projects
- Managed and successfully led BI and Maintenance projects for clients
- Engaged with clients like Genentech, SPSS, Coors, etc.
- Currently engaged as a Engagement Manager to manage Commercial SFA at Genentech, Inc.

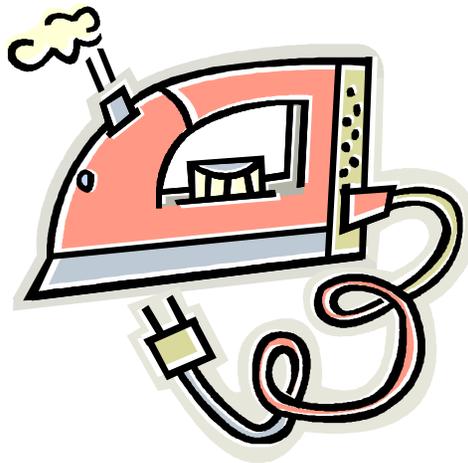
Evolution of Database Infrastructure



	1 st Generation	2 nd Generation
Features	<ul style="list-style-type: none"> • Standard database servers • Standard Storage (Local, SAN, NAS) • Standard processors (single core) • Serial processing • Processing in the database 	<ul style="list-style-type: none"> • Powerful database servers • Standard Storage (Local, SAN, NAS) • Improved processors (dual core, quad core) • Parallel processing • Processing in the database
Advantages / Disadvantages	<ul style="list-style-type: none"> • Slow processing • Index dependent • High maintenance • Multiple Vendors 	<ul style="list-style-type: none"> • Faster processors • Index dependent • Limitations on I/O • High maintenance • Multiple vendors
Vendors	<ul style="list-style-type: none"> • Oracle, Sybase, DB2 	<ul style="list-style-type: none"> • Oracle, Sybase, DB2

- Index dependent
- Limitation on I/O
- High maintenance
- Multiple vendors

What is an Appliance?



- **Out of the Box**
- **Low maintenance**
- **Performs a very specific function**

What is an Appliance?



BI Appliance

Hardware

- Servers
- Storage

Software

- Database
- OS
- Others



- Out of the Box
- Low maintenance
- Performs a very specific function
- High Performance

Why BI Appliance?



- **Large volume** data processing requirements
- Need for **greater speed** and throughput
- **Lower** development and maintenance **costs**
- **Lower dependency** on **multiple** H/W and S/W vendors

Evolution of Database Infrastructure



	1 st Generation	2 nd Generation	3 rd Generation	4 th Generation
Features	<ul style="list-style-type: none"> Standard database servers Standard Storage (Local, SAN, NAS) Standard processors (single core) Serial processing Processing in the database 	<ul style="list-style-type: none"> Powerful database servers Standard Storage (Local, SAN, NAS) Improved processors (dual core, quad core) Parallel processing Processing in the database 	<ul style="list-style-type: none"> Initial BI Appliances Standard Storage (Local, SAN, NAS) Improved processors (multi quad core) Massively Parallel processing Processing at Storage level and at database level 	<ul style="list-style-type: none"> Advanced BI Appliances Virtualization Storage in the Cloud (Private/Public) Infrastructure as a Service (IaaS) model Hadoop/MapReduce
Advantages / Disadvantages	<ul style="list-style-type: none"> Slow processing Index dependent High maintenance Multiple Vendors 	<ul style="list-style-type: none"> Faster processors Index dependent Limitations on I/O High maintenance Multiple vendors 	<ul style="list-style-type: none"> Black Box – less maintenance Fast processors Index independent No Limitations on I/O Infrastructure Heavy Sufficient for several terabytes 	<ul style="list-style-type: none"> Cloud – no maintenance Pay-as-you-go Fast processing Caters to Petabytes of data
Vendors	<ul style="list-style-type: none"> Oracle, Sybase, DB2 	<ul style="list-style-type: none"> Oracle, Sybase, DB2 	<ul style="list-style-type: none"> Teradata, Netezza, GreenPlum, Exadata 	<ul style="list-style-type: none"> GreenPlum (Private Cloud), Vertica (Public Cloud)

BI Appliance Implementation – A Case Study

Methodology adopted



Planning	Evaluation	Vendor Finalization	Implementation
<ul style="list-style-type: none">• Understand the Business Problem• Set Goals/Objectives• Setup a Governance Structure• Feasibility Study• Market Analysis• Vendor Demos• RFI to vendors• Vendor Shortlisting	<ul style="list-style-type: none">• Finalize POC Approach• Define Evaluation scenarios• Define Success Criteria• Setup Environment for POC• Invite Vendors for POC• Perform POC• Analyze POC Results• Finalize Vendor	<ul style="list-style-type: none">• Cost negotiations with Vendor• Funding Commit Process• Contract negotiations and finalization	<ul style="list-style-type: none">• Implementation Planning<ul style="list-style-type: none">- Data Center space- Network- Database migration- Power- Seismic isolation- Post-implementation support• Hardware Installation• Software Installation• Production Switchover• Post-Implementation Issues

Planning Phase



Planning	Evaluation	Vendor Finalization	Implementation
<ul style="list-style-type: none">• Understand the Business Problem• Set Goals/Objectives• Setup a Governance Structure• Feasibility Study• Market Analysis• Vendor Demos• RFI to vendors• Vendor Shortlisting	<ul style="list-style-type: none">• Finalize POC Approach• Define Evaluation scenarios• Define Success Criteria• Setup Environment for POC• Invite Vendors for POC• Perform POC• Analyze POC Results• Finalize Vendor	<ul style="list-style-type: none">• Cost negotiations with Vendor• Funding Commit Process• Contract negotiations and finalization	<ul style="list-style-type: none">• Implementation Planning<ul style="list-style-type: none">- Data Center space- Network- Database migration- Power- Seismic isolation- Post-implementation support• Hardware Installation• Software Installation• Production Switchover• Post-Implementation Issues

Business Problem



Planning

Evaluation

Vendor Finalization

Implementation

- Client undergoing a merger thereby increase in data volumes by 5X and user base by 3X
- Slow 'Time to Market' for new BI applications
- Query Performance Degradation
- Slow data loads

Goals/Objectives



Planning

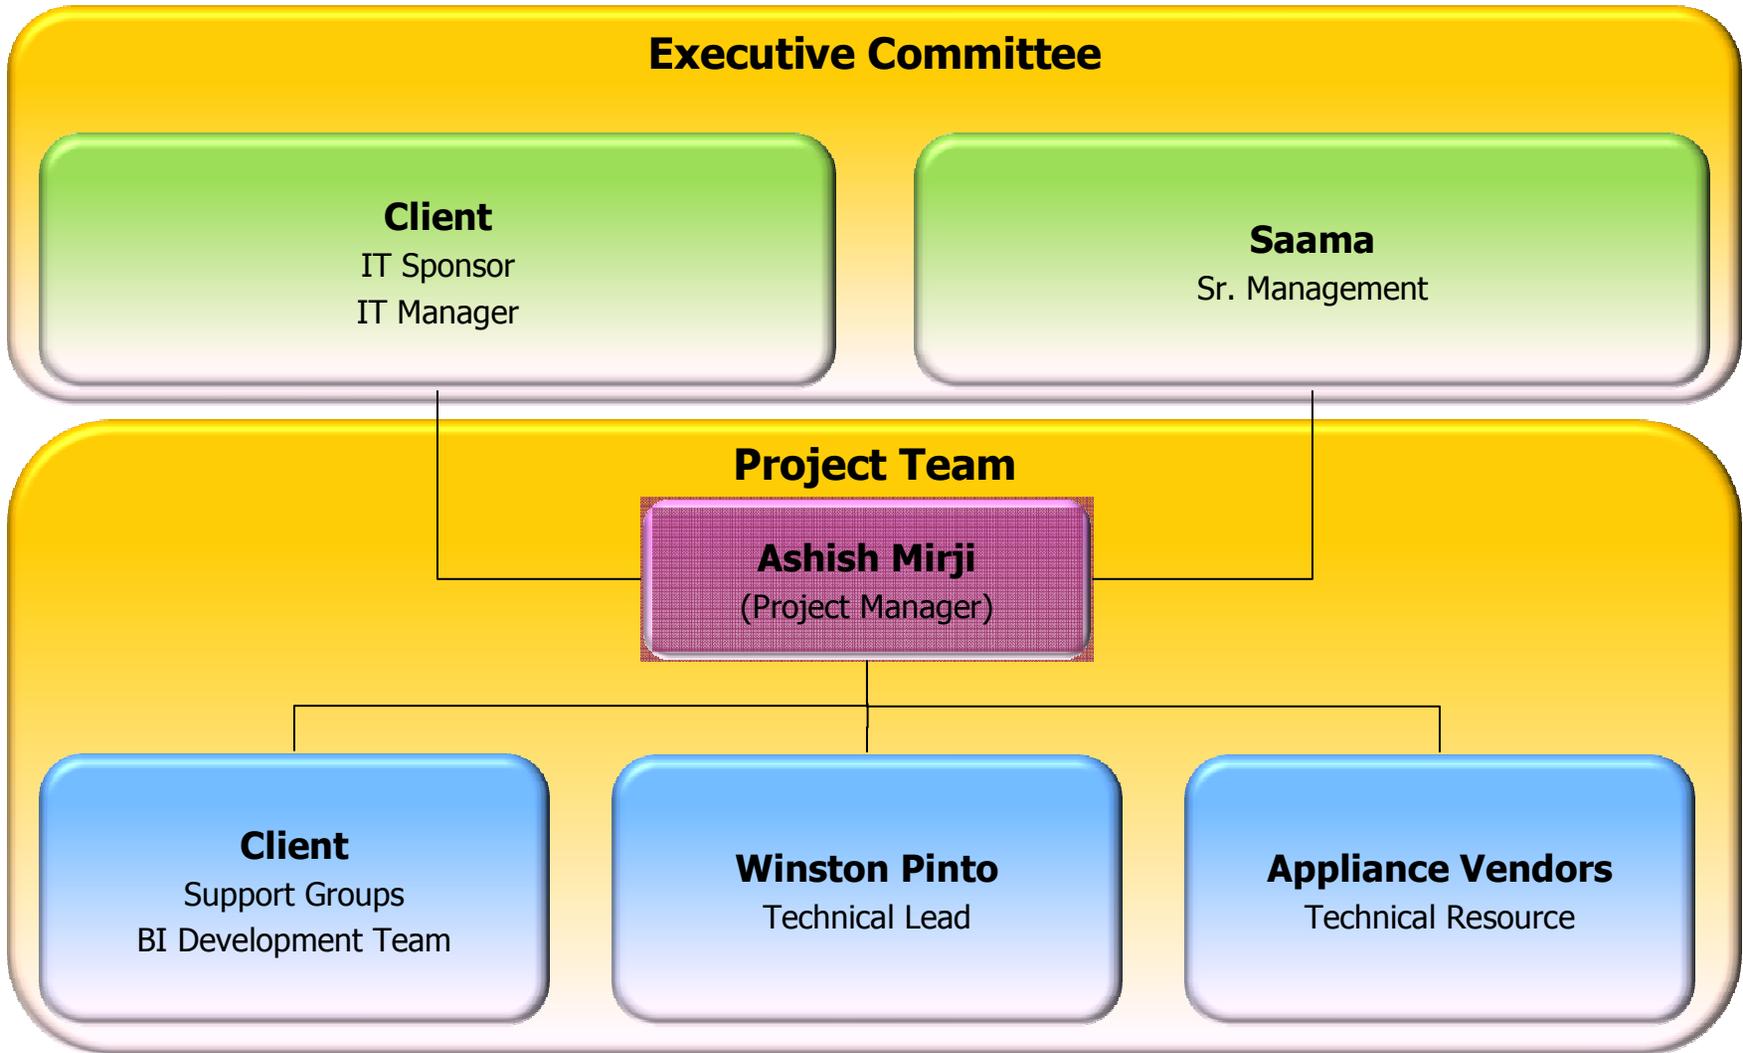
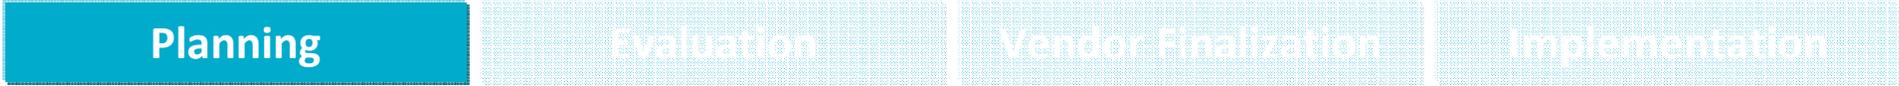
Evaluation

Vendor Finalization

Implementation

- **Improve 'Time to Market'** – Faster Development
- **Greater performance of queries** against atomic level data -
Reduction of additional objects currently used for aggregation
- **Faster Ad-hoc analysis** query performance
- **Faster Data Loads**
- **Ease of maintenance** and support
- **Ability to scale** to larger data volumes
- **Meet the SLAs** as promised to the Business
- **Lower development and M&E costs** (resources)

Governance Structure



Feasibility and Market Analysis



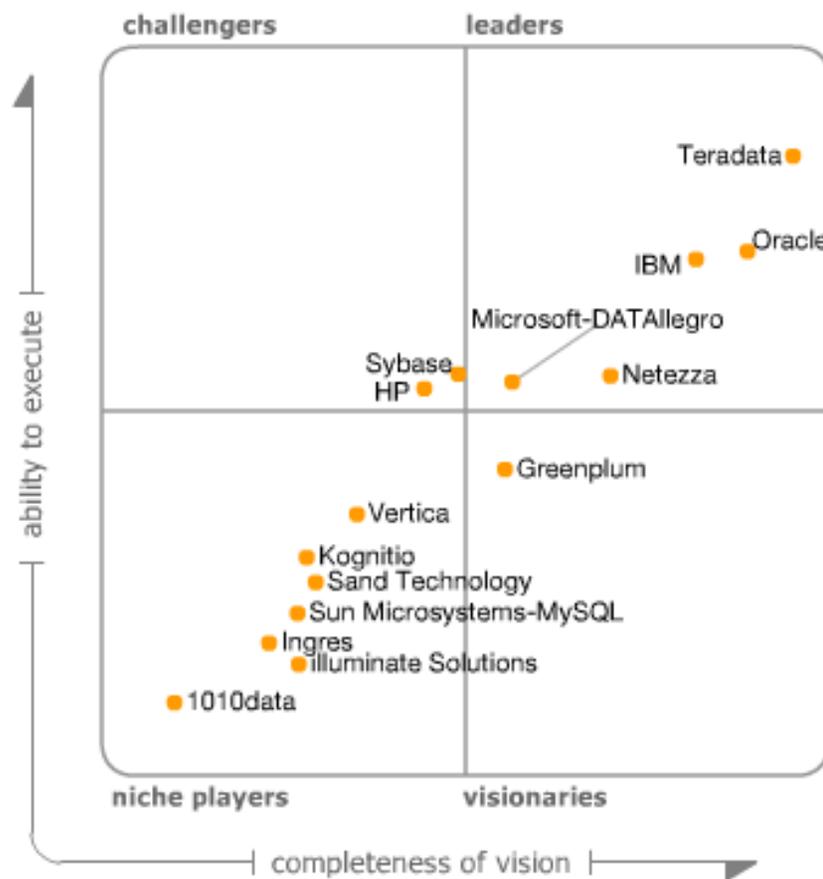
Planning

Evaluation

Vendor Finalization

Implementation

Gartner Magic Quadrant for BI Appliances 2009



www.gartner.com

Feasibility and Market Analysis



Planning

Validation

Vendor Finalization

Implementation

Types of BI Appliances available in Market

Software			Hybrid (H/W and S/W)
Traditional DB	Column-based DB	Accelerators	
GreenPlum	Vertica ParAccell Sybase	DataUpia SAP BIA	Netezza Teradata IBM Oracle Exadata

Feasibility and Market Analysis



Planning

Evaluation

Pre-Implementation

Implementation

Key selling features of BI Appliances

Massively Parallel Processing (MPP)

Columnar Compression

Apache Hadoop

MapReduce

Advanced data filtering - FPGA

Smart Flash Cache

Massively Parallel Processing (MPP)



Planning

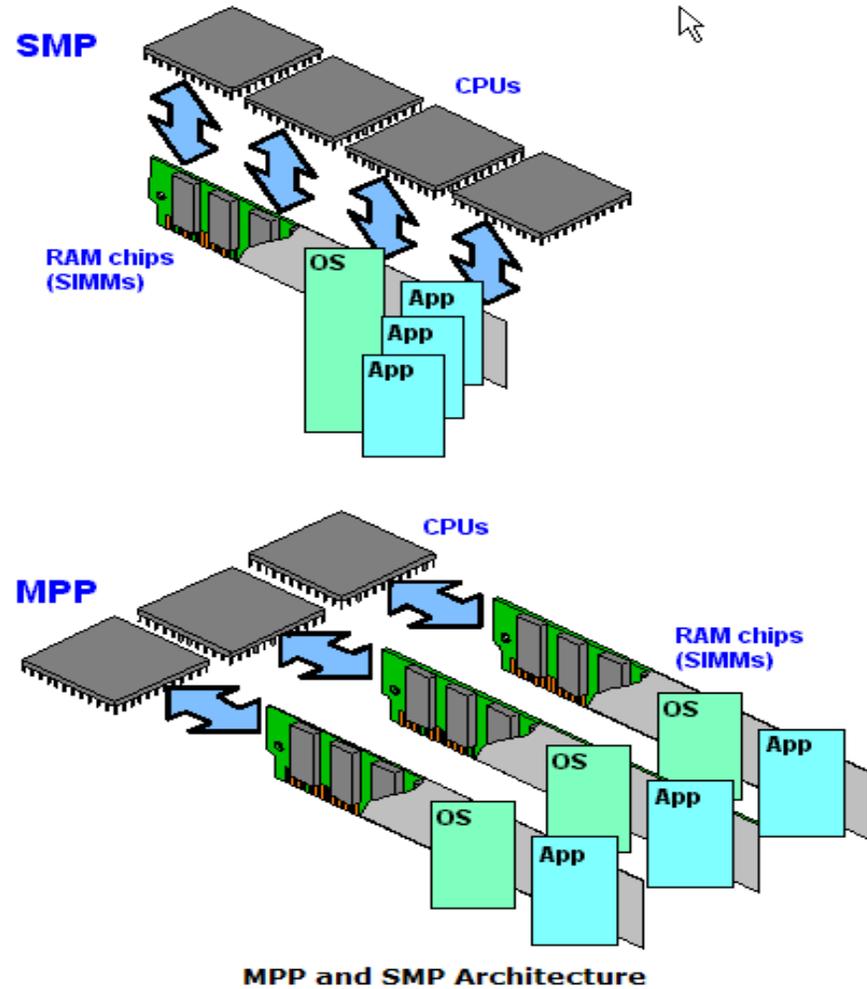
Evaluation

Pre-Implementation

Implementation

Massively Parallel Processing (MPP)

A multiprocessing architecture that uses many processors and a different programming paradigm than the common symmetric multiprocessing (SMP) found in today's computer systems



Hybrid Columnar Compression



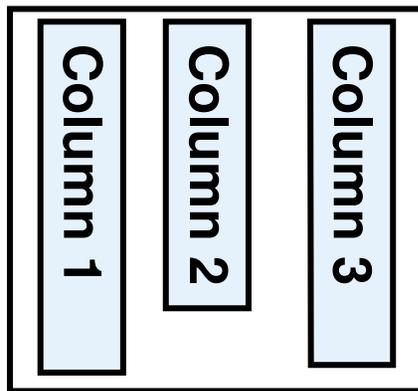
Planning

Evaluation

Pre-Implementation

Implementation

Compression Unit



- Tables are organized into sets of a few thousand rows called Compression Units (CUs)
- Within Compression Unit, data is Organized by Column and then compressed
 - Column organization brings similar values close together, enhancing compression
- Useful for data that is bulk loaded and queried
 - Update activity is light

Apache - Hadoop



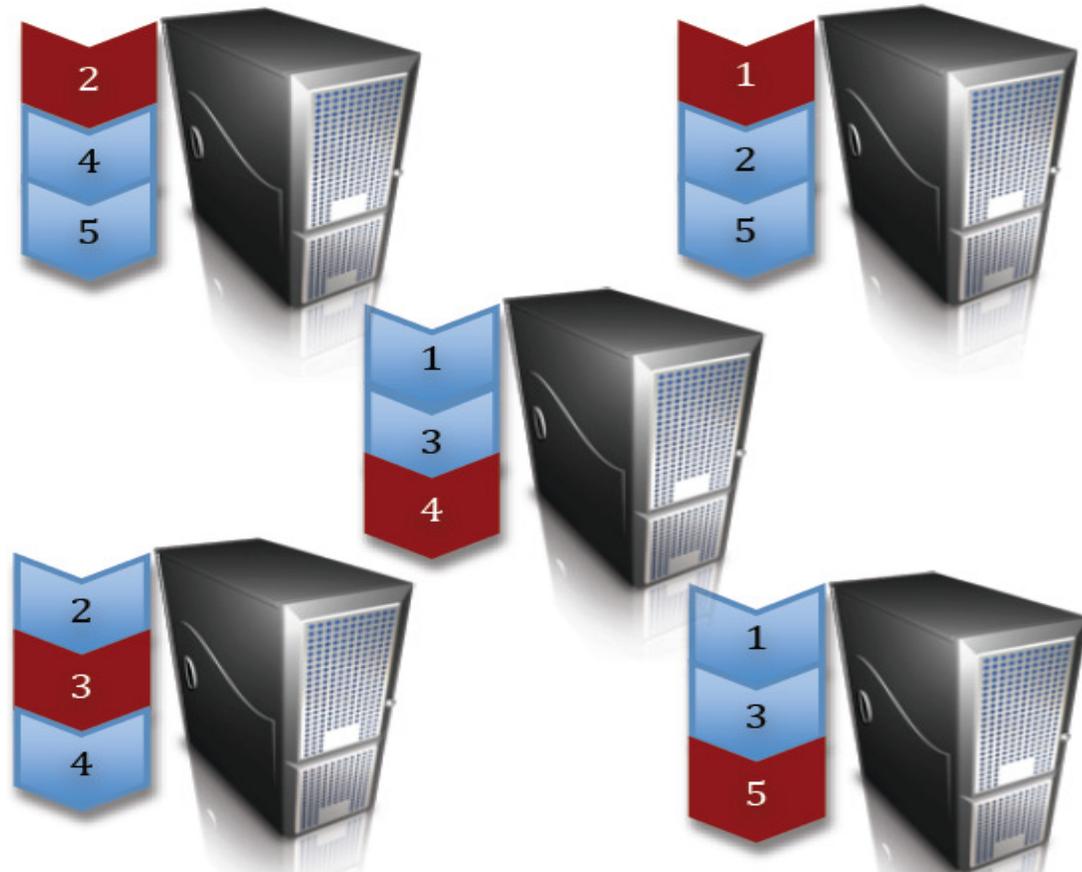
Planning

Evaluation

Pre-Implementation

Implementation

Hadoop takes advantage of HDFS' data distribution strategy to push work out to many nodes in a cluster. This allows analyses to run in parallel and eliminates the bottlenecks imposed by monolithic storage systems.



Map Reduce



Planning

Evaluation

Pre-Implementation

Implementation

MapReduce is a framework for processing huge datasets on certain kinds of distributable problems using a large number of computers (nodes), collectively referred to as a cluster. Computational processing can occur on data stored either in a [filesystem](#) (unstructured) or within a [database](#) (structured).

"Map" step: The master node takes the input, chops it up into smaller sub-problems, and distributes those to worker nodes. A worker node may do this again in turn, leading to a multi-level [tree](#) structure. The worker node processes that smaller problem, and passes the answer back to its master node.

"Reduce" step: The master node then takes the answers to all the sub-problems and combines them in a way to get the output - the answer to the problem it was originally trying to solve.

Advanced data filtering

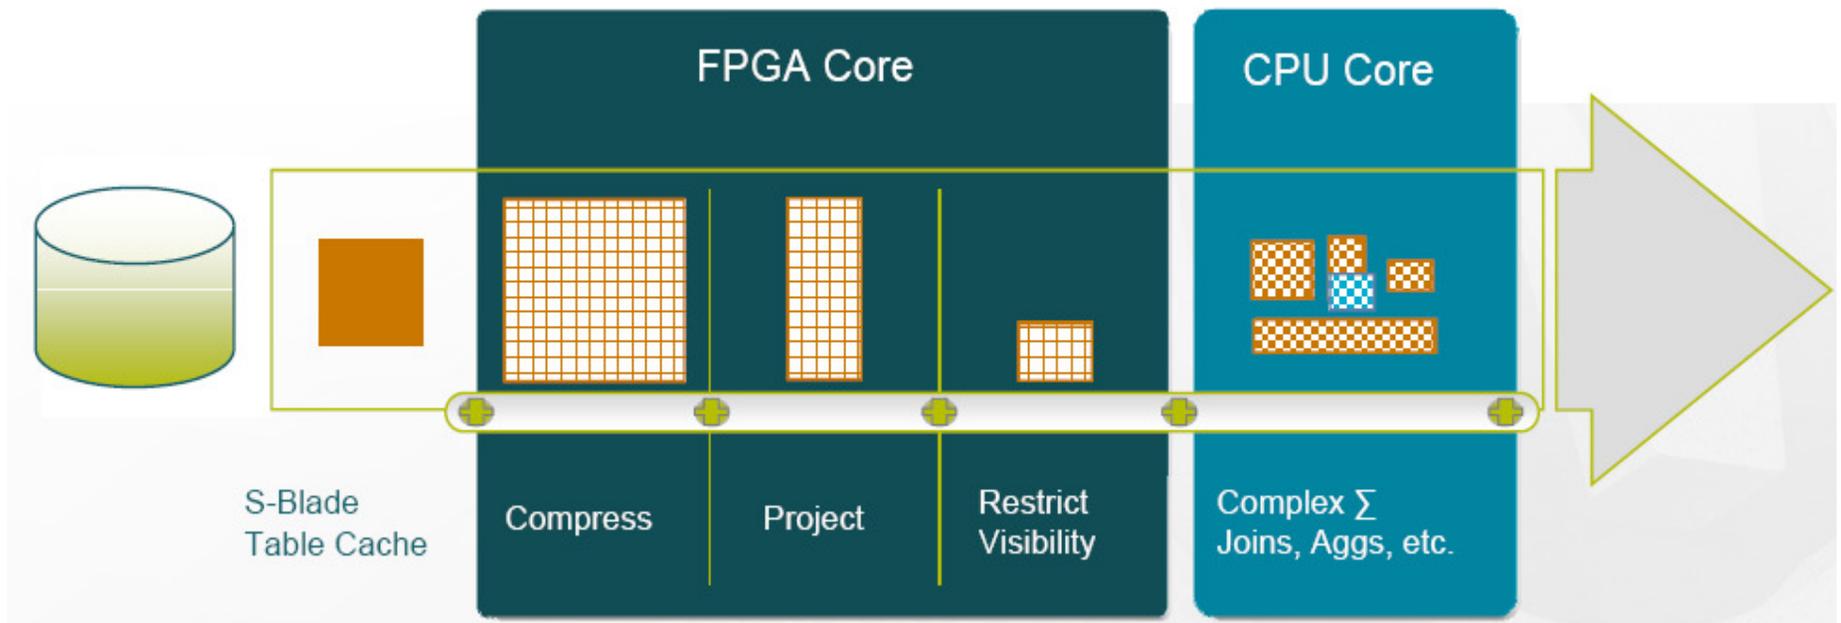


Planning

Evaluation

Pre-Implementation

Implementation



Smart Flash Cache



Planning

Evaluation

Pre-Implementation

Implementation

- Caches Hot Data Transparently in the Flash Cards
- Use PCI Express based Flash Cards for greater throughput and IOPs and avoid disk controller limitations
- Smart Caching
 - Smarter than basic LRU algorithm
 - Knows when to skip caching objects to avoid polluting or flushing the cache
- Allows applications to explicitly optimize caching



RFI Categories



Planning

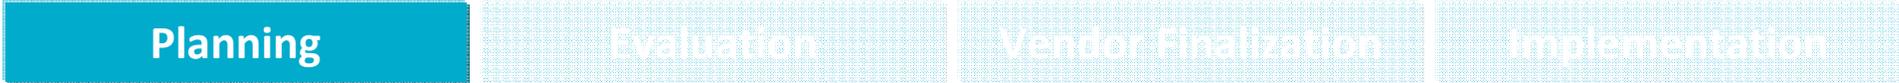
Evaluation

Vendor Finalization

Implementation

- **Vendor Background** – Company history, Financial stability, future product plans, etc.
- **Proposed Solution** – vendor’s recommended solution with reasons
- **Technical Support** – Cost, levels of support, coverage, SLAs, scope, etc.
- **Product Documentation** – Online, CD, Books, etc.
- **Training** – Classes, certifications, cost, etc.
- **Partnerships** – Partnerships with other BI/DW market vendors
- **Technical Architecture** – architecture components, roadmap, SaaS model, etc.
- **Integration with other BI/DW tools in market** – APIs to connect to other tools
- **Platforms supported** – H/W, OS supported etc.
- **Scalability/Development/Administration** – Ease of development and administration
- **Customer References** – Case study, endorsements, phone calls, relevant to industry
- **Pricing** – H/W, S/W, support, upgrades, perpetual/subscription model

Shortlisting process



Non-Functional Requirements Matrix

#	Criteria	Critical	Weight Factor (1-10)	Weight Factor (%)	IBM	Netezza	Greenplum	Teradata	Sybase	Dataupia	Oracle	SAP BIA	ParAccell
Non-Functional Requirements													
Vendor Background													
1	Vendor has a strong market presence in BI		5	26.32%									
2	Vendor is financially strong	TRUE	8	42.11%									
3	Vendor has a good customer base across verticals and geographies		6	31.58%									
	Total		19	100.00%									
Proposed Solution Overview													
4	Proposed solution has been implemented at many customers		8	30.77%									
5	Proposed solution has been implemented at many customers in health sciences	TRUE	10	38.46%									
6	Proposed solution has a strong market share in the BI Apps space		8	30.77%									
	Total		26	100.00%									
Technical Support & Training													
7	Strong technical support	TRUE	10	30.30%									
8	Online User forum, user group etc for exchange of knowledge		7	21.21%									
9	Online knowledge base of previously resolved issues		5	15.15%									
10	Training classes available	TRUE	8	24.24%									
11	CBT and/or other offline training options		3	9.09%									
	Total		33	100.00%									
Implementation & Strategic Partnerships													
12	Duration of implementation cycle		8	30.77%									

Shortlisting process



Planning

Evaluation

Vendor Finalization

Implementation

Functional Requirements Matrix

#	Criteria	Critical	Weight Factor (1-10)	Weight Factor (%)	IBM	Netezza	Greenplum	Teradata	Sybase	Dataupia	Oracle	SAP BIA	ParAccell
Functional Requirements													
Integration													
1	Interface with Informatica (Native Connectivity, bulk load support)	TRUE	8	23.53%									
2	Interface with Business Objects (Universe, Crystal, WEBI, Deskys)	TRUE	8	23.53%									
3	Metadata exchange with ETL and reporting tools (using Meta Integration)		4	11.76%									
4	Interface with Modeling tools (ER Studio)		5	14.71%									
5	Support for the following interfaces – ODBC, JDBC, SQL, .NET, J2EE, XML, Web		6	17.65%									
6	Tools to migrate apps from Oracle to your solution		3	8.82%									
	Total		34	100.00%									
Platform Support													
7	Hardware support - is it generic? If not, does it specifically exploit the features of your		5	20.00%									
8	OS support - is it generic? If not, does it specifically exploit the features of your		5	20.00%									
9	Storage support - is it generic? If not, does it specifically exploit the features of your		5	20.00%									
10	Platforms for client tools (PC, Mac)		5	20.00%									
11	Platforms for web-based client tools (Safari, Firefox, Internet Explorer)		5	20.00%									
	Total		25	100.00%									
Performance													
12	Query optimization for complex queries	TRUE	10	28.57%									
13	User and/or query prioritization		5	14.29%									
14	Automatic aggregate build		5	14.29%									
15	Minimal effort required to tune queries	TRUE	10	28.57%									
16	Data partitioning, indexing or similar strategy for better query response		5	14.29%									

Shortlisted Candidates for further evaluation



Planning

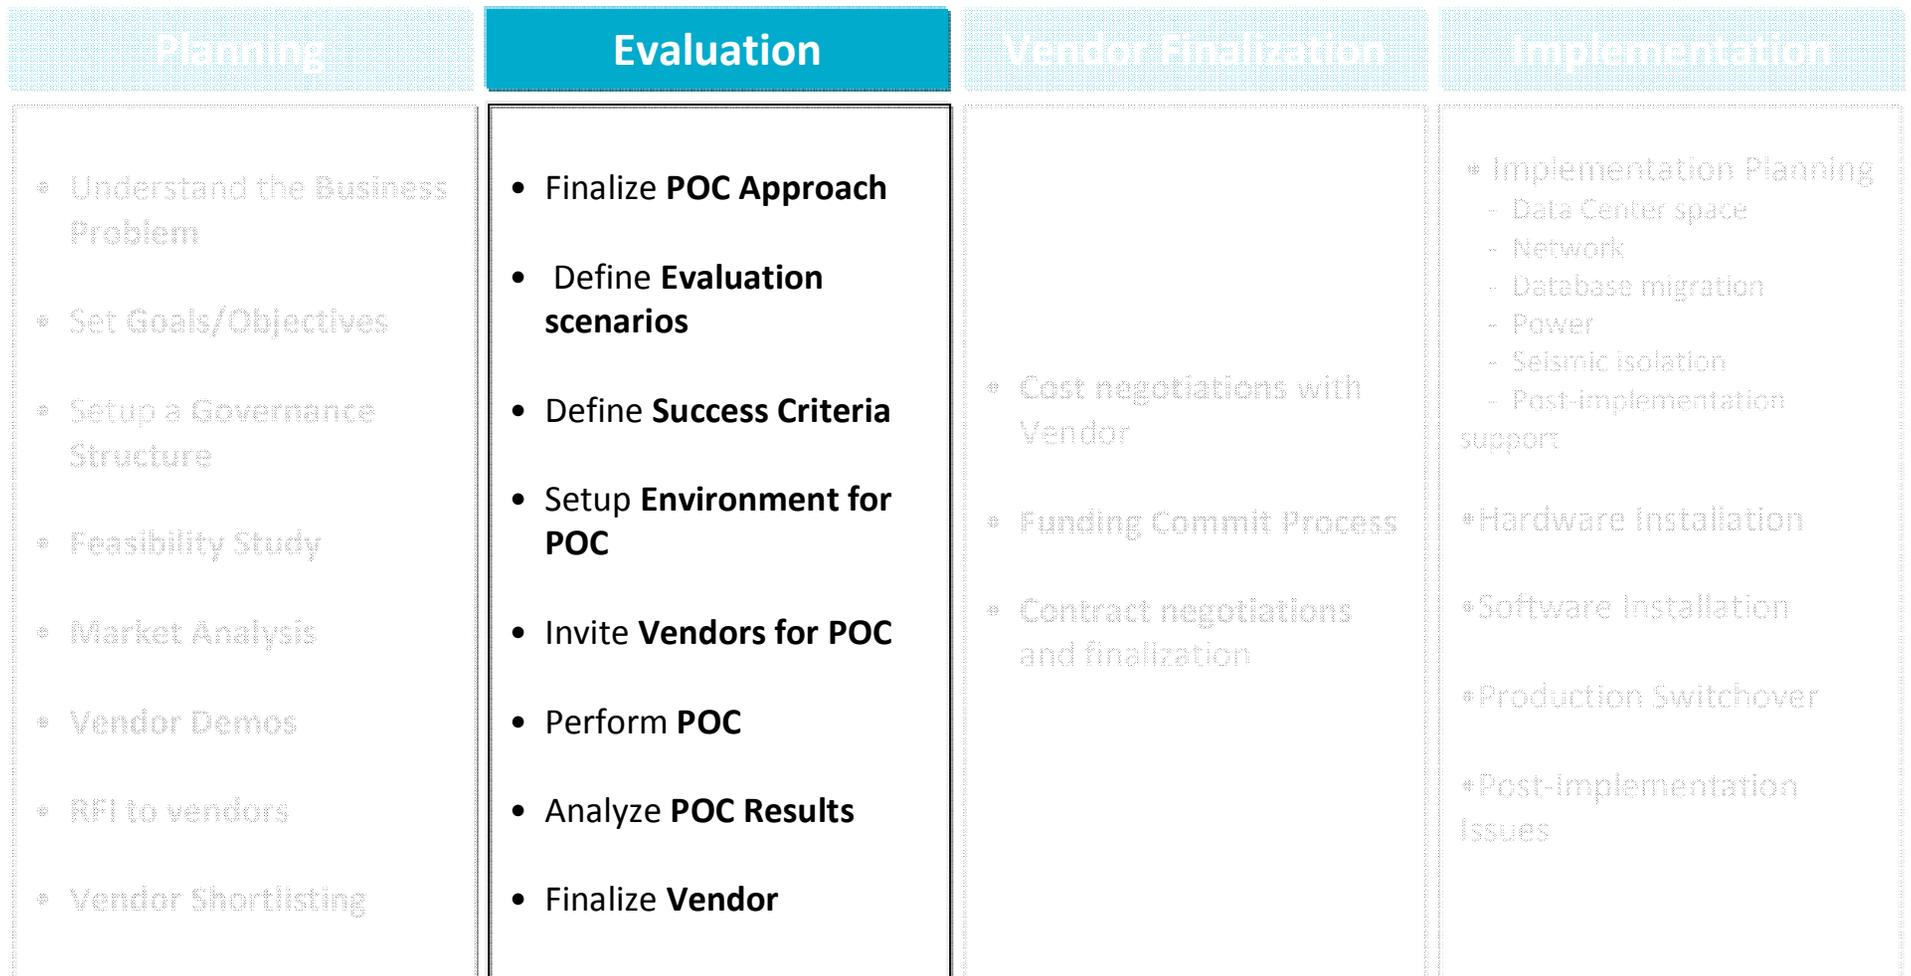
Evaluation

Vendor Evaluation

Implementation

Criteria	Vendor 1	Vendor 2	Vendor 3	Vendor 4
Type	Hybrid	Software	Hybrid	Hybrid
Selling Feature	MPP, Hybrid Columnar Compression, Smart Flash Cache	MPP, MapReduce, Virtualization	MPP, FPGA	MPP
Reason shortlisted	Best of breeds, easy conversion, existing relationship, possible low TCO, financially stable	Flexibility to use any hardware, possible lowest TCO, emerging company, good clientele, financially stable	Possible low TCO, good clientele, pioneers in low cost BI Appliances, financially stable	Oldest vendor in this space, good clientele, financially stable, new low cost solutions available

Evaluation Phase



Sample Evaluation Scenarios (ETL)



Planning

Evaluation

Vendor Finalization

Implementation

ETL

Sr. No.	Factor	Scenario	What does this test?
1.	Performance	Run the Informatica workflow to load the current daily and monthly file against the BI Appliance	Compares the performance between the existing Oracle DB versus the BI Appliance
2.	Scalability	Run the Informatica workflow to load the daily and monthly file with data 10X to 100X times of normal against the BI Appliance	Tests the scalability and how the BI Appliance will perform with increased loads
3.	Compatibility	Connect to Informatica. Check to see if the connection is via a 'native driver' or an ODBC	Tests the compatibility of the BI Appliance to Informatica
4.	DB Feature	Run Informatica workflows that insert/delete/update the data in the database	Checks to see similar database operations on the BI Appliance.
5.	INFA/DB Feature	Run the Informatica workflow and break the session to see if the data is rolled back	Checks to see if the BI Appliance rollbacks the changes made
6.	Migration	Run the Informatica workflow that contains Oracle Hints, Functions or Stored Procedures.	Checks to see how the DB features like Hints, SPs, Functions are handled within the BI Appliance and what is the effort taken to re-write the code.
7.	DB Feature	Run the Informatica workflow that manages the partition rollover (delete partition/create partition)	Checks to see how the BI Appliance supports this scenario

Sample Evaluation Scenarios (Query/Reporting)



Planning

Evaluation

Vendor Finalization

Implementation

Query/Reporting

Sr. No.	Factor	Scenario	What does this test?
1.	Performance	Run current long-running queries against aggregates using the Appliance	Checks performance as compared to that in the current Oracle environment
2.	Performance	Run current long-running queries against the base FACT tables	Confirms the fact that Appliances can be faster even against the base FACT tables and there is no need for aggregates
3.	Performance	Run the list of queries sequentially	Check the Checks performance as compared to that in the current Oracle environment
4.	Performance	Run the list of queries in parallel (atleast 20 concurrent transactions)	Check the Checks performance as compared to that in the current Oracle environment
5.	Scalability	Run the list of queries in parallel (atleast 40 concurrent transactions)	Check that the performance when the number of concurrent transactions are doubled.
6.	Performance	Run the Informatica workflow to load X number of tables and simultaneously query the same tables in multiple concurrent sessions	Check the effect of querying the tables while they are being loaded on the performance of loads and queries
7.	Scalability	Run the same queries (above) against 10X to 100X times the underlying data	Checks that the queries perform satisfactorily after scaling up the data
8.	Compatibility	Connect Business Objects to the Appliance and run basic queries/tasks to check the compatibility	Check to see the type of connectivity (native driver versus ODBC/JDBC)
9.	Migration	Run queries or reports that have specific syntax used by Oracle and note the effort taken to re-write the code	Checks the effort to migrate the code to the Appliance
10.	Migration	Check to see the availability of equivalent for Oracle Functions for aggregation, multi-period analysis (LEAD, LAG, etc.) and Oracle Hints	Checks the effort to migrate the code to the

Sample Evaluation Scenarios (Database/Data Model)



Planning

Evaluation

Vendor Finalization

Implementation

Database

Sr. No.	Factor	Scenario	What does this test?
1.	Migration	Check availability of DB objects like Stored procedures, Constraints, Partitions, etc. and measure the effort taken to convert them to the equivalent in the BI Appliance	Checks the effort required for the migration from Oracle to the BI Appliance
2.	Administration	Take a backup of the database and recover the same	Checks the availability of DB features like Backup and Recovery and the process for the same
3.	Administration	Understand the process of applying Patches and performing Upgrades	Checks to see the available tools and the process of applying patches and upgrades.

Data Model

Sr. No.	Factor	Scenario	What does this test?
1.	Architecture	If possible, migrate a 3NF model to the appliance	Check whether the appliance works for a 3NF model as well as a Denormalized model
2.	Migration	Run scripts that depend on the database referential integrity to see how they perform in the BI Appliance and understand the effort required to modify the code (if required)	Checks the availability of DB features like constraints (Null, PK, FK, etc.) and effort required to migrating these features
3.	Compatibility	Use the modeling tool (used at Gene) to connect to the BI Appliance DB. Check to see if scripts can be generated using the modeling tool.	Checks the compatibility of the Appliance with the modeling tools

Sample Evaluation Scenarios (Server/Hardware/Tools)



Planning

Evaluation

Vendor Finalization

Implementation

Server/Hardware

Sr. No.	Factor	Scenario	What does this test?
1.	Performance	Record the CPU, memory usage and IO with current ETL/Reports	Compares the performance with that of existing environment
2.	Scalability	Record the CPU, memory usage and IO with 10X to 100X times data for ETL/Reports	Checks the load on the server due to increased load
3.	Scalability	Increase the capacity of the Server and run the same loads (ETLs) and queries (Reports)	Check if the performance linearly scales as the server capacity is increased
4.	Compatibility	Check if external storage devices (SAN/NAS) can be used and whether the system can be on a virtual machine	Checks the compatibility of the server with features relevant to standard database servers
5.	Administration	Understand the DR/Failover mechanism within the Appliance (Node Failure/Disk Failure)	Helps us understand the administration feature specific to failover/Disaster Recovery

Tools/Utilities

Sr. No.	Factor	Scenario
1.	Development	Check if the appliance has any tools for development similar to that available with Oracle (PLSQL developer, SQL Loader, etc.)
2.	Administration	Check if the appliance has any tools for administration similar to that available with Oracle (User administration, tuning, debugging, space management , RMAN, etc.)
3.	Monitoring	Check if the appliance has any tools for monitoring similar to that available with Oracle (Statspack, etc.)

Selection Criteria



Planning

Evaluation

Vendor Finalization

Implementation

- The implementation of the BI Appliance should **significantly lower the Total Cost of Ownership (TCO)** for the company
- The implementation of the BI Appliance should significantly **lower the development time** and hence the 'Time to Value' to the Business
- BI Appliance should **improve the performance** of the queries, reports and ETL processes by at least 10X times
- There should be **minimal effort required for implementation** of the BI Appliance and **minimal effort for migration** of existing data from the current infrastructure to the Appliance Database
- There should be **minimal support and administrative overhead** required for ongoing support effort of the BI Appliance
- BI Appliance should have all the required **tools and methods** required for development and basic database and server administration
- BI Appliance should be **compatible** with the existing Data Integration, Reporting and EII tools
- BI Appliance should be linearly **scalable** to accommodate 200+ concurrent users and 20+ TB of user data
- BI Appliance should have the required features and tools to maintain **security** of data

Vendors selected for POC



Planning

Evaluation

Vendor Finalization

Implementation

- Two vendors were eventually selected for the POC after going through several discussions with the vendors and evaluating the products against various parameters
- **Vendor 1**
 - Adopted the MPP architecture similar to Teradata and Netezza which has proved very effective with their customers
 - Incumbent vendor
 - Lowest TCO
 - Leverage Database and RAC Licenses
 - Skills readily available
 - Minimal implementation & conversion cost
- **Vendor 2**
 - Ability to use commodity hardware (Netezza and Teradata will not evolve as fast in taking advantage of newer hardware)
 - Ability to virtualize (conducting a POC with VMWare)
 - Licensing by 'User Data Volume' rather than CPUs
 - Lower TCO than Netezza and Teradata

Pre-POC Preparation (Benchmarking)



Planning

Evaluation

Vendor Selection

Implementation

- Define the scope for Benchmarking – ETLs, Queries, Reports
- Identify the environment for Benchmarking
- Refresh database with production data
- Connect ETL and Reporting tools to the new environment
- Configure QA tool (Loadrunner)
- Execute Test Cases (3 iterations each)
 - Scenarios –
 - Run queries sequentially as a single user
 - Run queries sequentially for 25 concurrent users with a lag of 2 minutes
 - Increase the size of the data to 5X and repeat the above scenarios
- Define success criteria (expected run timing) for each query

POC Process



Planning

Evaluation

Vendor Finalization

Implementation

- Rules
 - The FACT tables will be partitioned and sub-partitioned as per the current environment
 - No indexes or hints to be used
 - Monitor the CPU, Memory and I/O Usage on each of the nodes at regular intervals (~ 2 minutes) for the single user and concurrency tests
- Steps
 - ¼ Rack test
 - Import data without indexes into the database and expand data to 4TB
 - Execute queries for single user and 25 concurrent users
 - Insert data from flat file into the database and check for compression
 - ½ Rack test
 - Redistribute the same 4TB data to ½ Rack and run similar tests

POC Results



Planning

Evaluation

Vendor Finalization

Implementation

Test	Query	Current Env. (seconds)	Success Criteria (seconds)	Vendor 1 Results (seconds)	Vendor 2 Results (seconds)	Comments
4TB – Single User Test	1	931	45	2.7	27	Vendor 1 out performed Vendor 2 and exceeded the success criteria.
	10	237	9	0.9	1	
	15	1386	69	6.3	794	
4TB – Concurrent User Test	1	1273	35	16.3	117	Vendor 1 exceeded the expectations. Vendor 2 did not perform well on all the concurrency tests
	10	337	8	7.5	112	
	15	785	27	24.2	22	

Vendor Finalization Phase



Planning	Evaluation	Vendor Finalization	Implementation
<ul style="list-style-type: none"> • Understand the Business Problem • Set Goals/Objectives • Setup a Governance Structure • Feasibility Study • Market Analysis • Vendor Demos • RFI to vendors • Vendor Shortlisting 	<ul style="list-style-type: none"> • Finalize POC Approach • Define Evaluation scenarios • Define Success Criteria • Setup Environment for POC • Invite Vendors for POC • Perform POC • Analyze POC Results • Finalize Vendor 	<ul style="list-style-type: none"> • Cost negotiations with Vendor • Funding Commit Process • Contract negotiations and finalization 	<ul style="list-style-type: none"> • Implementation Planning <ul style="list-style-type: none"> - Data Center space - Network - Database migration - Power - Seismic isolation - Post implementation support • Hardware Installation • Software Installation • Production Switchover • Post-Implementation Issues

Cost Negotiations – Things to consider



Planning

Evaluation

Vendor Finalization

Implementation

- Hardware Cost – Servers/Racks
- Database Software Cost – DB, Mgmt and Productivity tools
- Appliance Software Cost – Storage S/W
- Increments each year – 5% over previous year
- Warranty – 3 yr for H/W and 1 yr S/W
- Paid Support – levels of support
- Upgrades – H/W, OS, firmware, DB, etc.

TCO Comparison and Vendor Finalization



Planning

Evaluation

Vendor Finalization

Implementation

Current Environment						
<u>Cost Breakdown</u>	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Project Costs (for Performance Tuning only)						
Current projects (15% of 2010 budget)	\$ 625,000.00					\$ 625,000.00
Costs due to merger (15% of 2010 budget)	\$ 240,000.00					\$ 240,000.00
Other Projects (approx)		\$ 250,000.00	\$ 250,000.00	\$ 250,000.00	\$ 250,000.00	\$ 1,000,000.00
M&E (10% of 2010 budget)	\$ 96,000.00	\$ 96,000.00	\$ 96,000.00	\$ 96,000.00	\$ 96,000.00	\$ 480,000.00
Hardware (\$5/GB/Month)	\$ 360,000.00				\$ 360,000.00	\$ 720,000.00
Total	\$ 1,321,000.00	\$ 346,000.00	\$ 346,000.00	\$ 346,000.00	\$ 706,000.00	\$ 3,065,000.00
Vendor 1 (Perpetual Model)						
<u>Cost Breakdown</u>	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Hardware Cost	\$ 525,000.00	\$ -	\$ -	\$ -	\$ -	\$ 525,000.00
Appliance Software Cost	\$ 129,600.00					\$ 129,600.00
Other Software Cost	\$ 177,120.00					\$ 177,120.00
Software Support Cost	\$ 35,798.00	\$ 37,229.92	\$ 38,719.12	\$ 40,655.07	\$ 42,687.83	\$ 195,089.94
Hardware Support Cost	-	-	-	\$ 23,000.00	\$ 23,000.00	\$ 46,000.00
Other Support Cost	\$ 31,680.00	\$ 32,947.20	\$ 34,265.09	\$ 35,978.34	\$ 37,777.26	\$ 172,647.89
Implementation Cost	\$ 50,000.00					\$ 50,000.00
Project Costs (for Performance Tuning only)						
Current projects (25% of Current Environment)	\$ 131,250.00					\$ 131,250.00
Costs due to Merger (25% of Current Environment)	\$ 8,949.50					\$ 8,949.50
Other Projects						\$ -
M&E (25% of Current Environment)	\$ 24,000.00					\$ 24,000.00
Support FTE	\$ -	\$ -	\$ -	\$ -	\$ -	\$ -
Total	\$ 1,113,397.50	\$ 70,177.12	\$ 72,984.20	\$ 99,633.42	\$ 103,465.09	\$ 1,459,657.33
Vendor 2 (Subscription Model)						
<u>Cost Breakdown</u>	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Appliance Cost	\$ 485,000.00	\$ 209,000.00	\$ 209,000.00	\$ 125,000.00	\$ 213,000.00	\$ 1,241,000.00
Support Cost	\$ -	\$ 30,000.00	\$ 50,000.00	\$ 75,000.00	\$ 90,000.00	\$ 245,000.00
Implementation Cost	\$ 300,000.00					\$ 300,000.00
Project Costs (for Performance Tuning only)						
Current projects (25% of Current Environment)	\$ 131,250.00					\$ 131,250.00
Costs due to Merger (25% of Current Environment)	\$ 8,949.50					\$ 8,949.50
Other Projects						\$ -
M&E (25% of Current Environment)	\$ 24,000.00					\$ 24,000.00
Support FTE (1 number)						\$ -
Total	\$ 949,199.50	\$ 239,000.00	\$ 259,000.00	\$ 200,000.00	\$ 303,000.00	\$ 1,950,199.50

Funding



Planning

Evaluation

Vendor Finalization

Implementation

- Architecture Review Board (ARB) approval
 - Fits within the Enterprise Architecture
- Capital Spending Request
 - Justification for spend
 - Capital vs. Expense
 - Over 5 years
- Funding approval

Contract Negotiation



Planning

Evaluation

Vendor Finalization

Implementation

- Contract agreement documents from vendor
- Management Review
- Legal Review
- Contract Negotiation

Implementation Phase



Planning	Evaluation	Vendor Finalization	Implementation
<ul style="list-style-type: none">• Understand the Business Problem• Set Goals/Objectives• Setup a Governance Structure• Feasibility Study• Market Analysis• Vendor Demos• RFI to vendors• Vendor Shortlisting	<ul style="list-style-type: none">• Finalize POC Approach• Define Evaluation scenarios• Define Success Criteria• Setup Environment for POC• Invite Vendors for POC• Perform POC• Analyze POC Results• Finalize Vendor	<ul style="list-style-type: none">• Cost negotiations with Vendor• Funding Commit Process• Contract negotiations and finalization	<ul style="list-style-type: none">• Implementation Planning<ul style="list-style-type: none">- Data Center space- Network- Database migration- Power- Seismic isolation- Post-implementation support• Hardware Installation• Software Installation• Production Switchover• Post-Implementation Issues

Implementation Planning



Planning

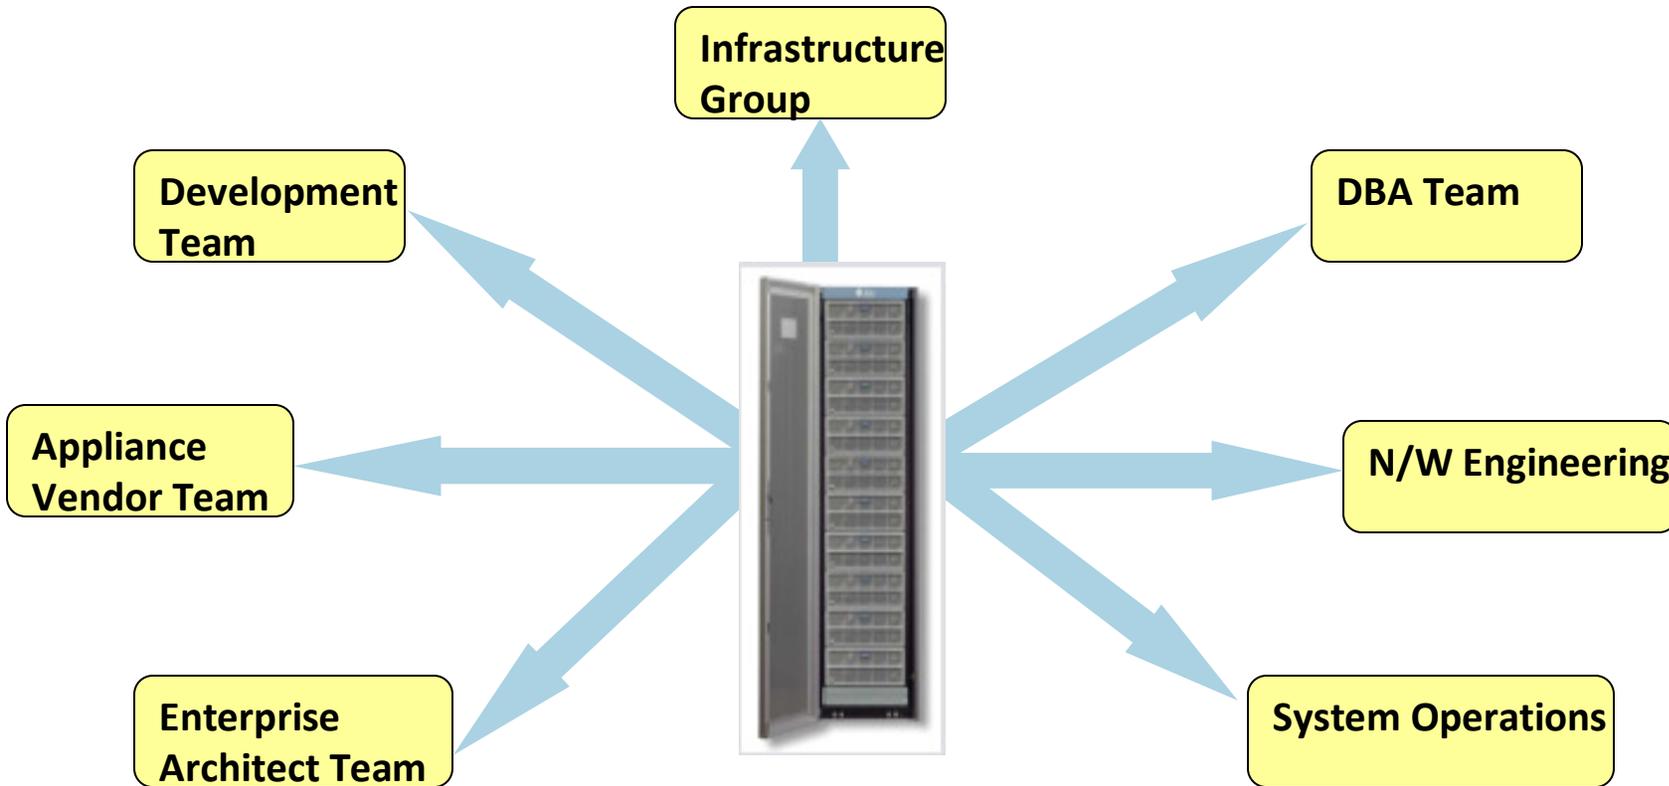
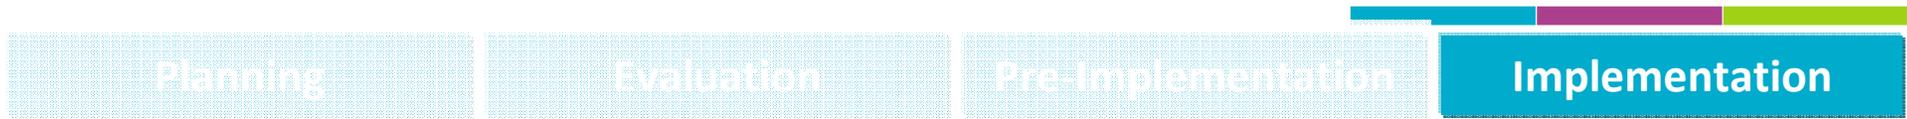
Evaluation

Pre-Implementation

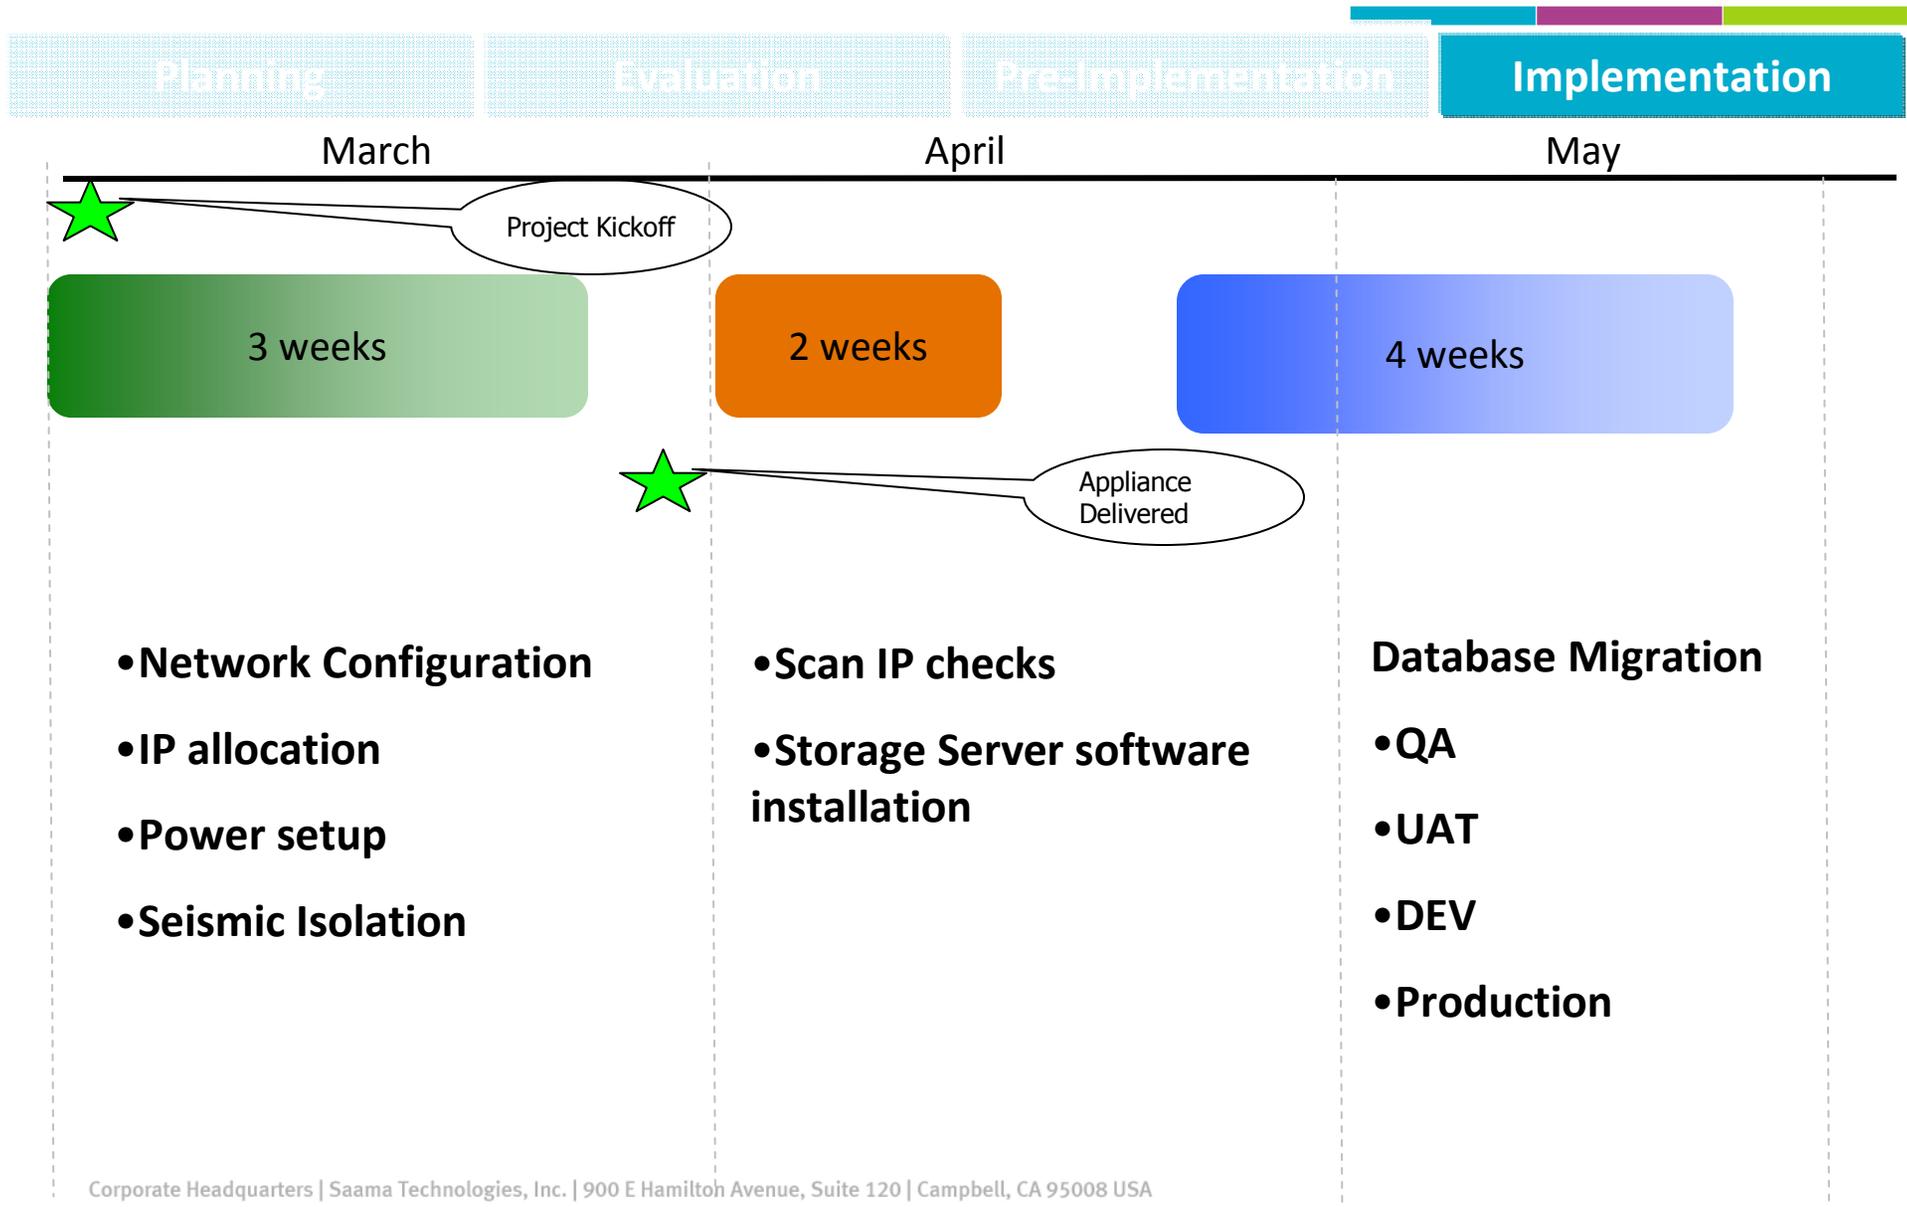
Implementation

- Data Center space
- Network
- Database migration
- Power
- Seismic isolation
- Post-implementation support

Team Involvement



Schedule



Network Planning for Quarter Rack

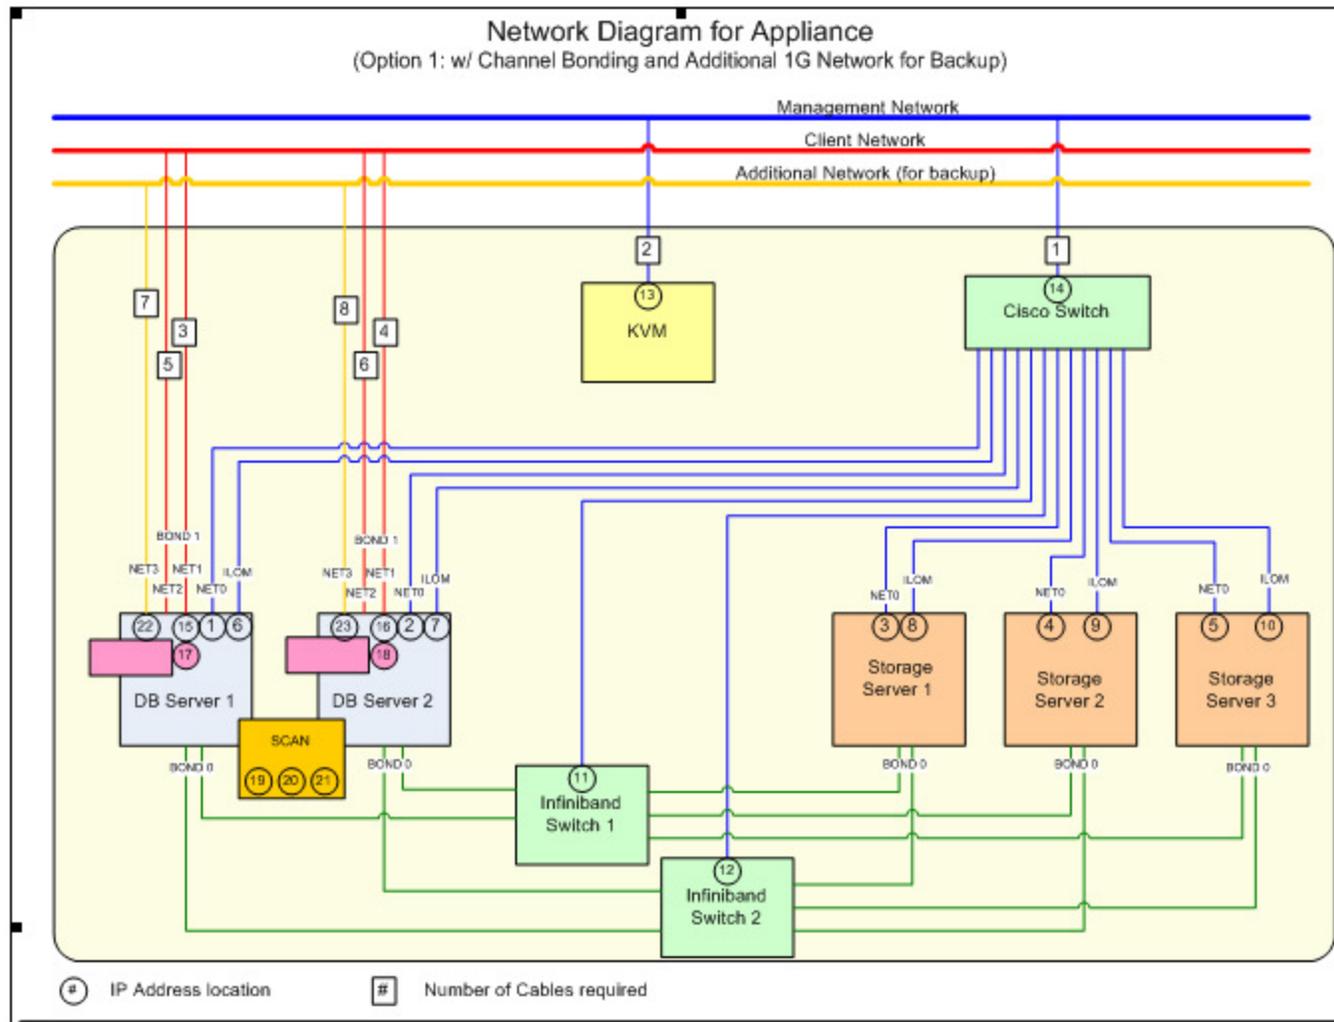


Planning

Evaluation

Pre-Implementation

Implementation



Migration Strategy

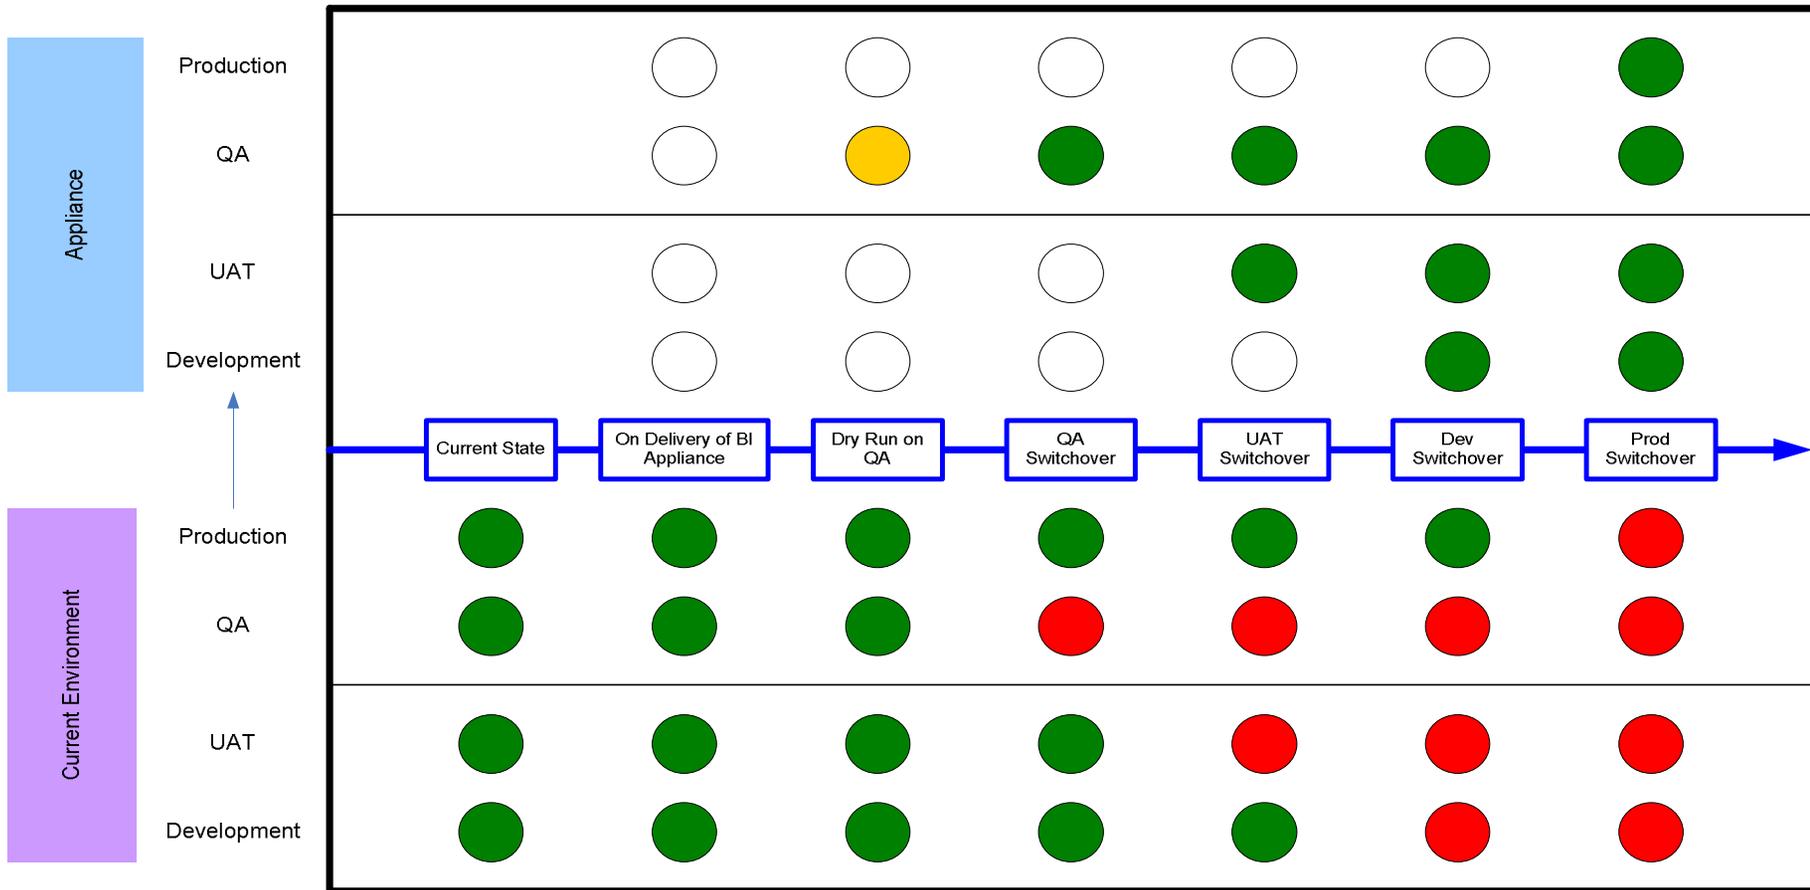


Planning

Evaluation

Pre-Implementation

Implementation



Migration Best Practices



Planning

Evaluation

Pre-Implementation

Implementation

- Configure the Allocation Unit Size to 4 MB
- Configure Optimal Database Extent Sizes
- Choose the Right Migration Strategy

Post Migration Checks



Planning

Evaluation

Pre-Implementation

Implementation

- Check Disk Groups for Rebalance
- Access Index Requirements

Performance Benefits



Planning

Evaluation

Pre-Implementation

Implementation

	Original	Actual	% Gain
INFA	12:40:00	1:30:00	80 %
BOBJ	Overall Average Gain of 80 %		
Stored Procedure	3:30:00	0:30:00	86%
Data Archival	1:05:00	0:22:00	66%
Index Creation	0:04:52	0:01:49	63 %
Table Stats	1:19:41	0:02:39	97%

Challenges Faced



Planning

Analysis

Pre-Implementation

Implementation

- Coordination with various teams
- Project Go Live dates clashing with Appliance migration
- Least impact to Development team and Business

Post Implementation Issues



Planning

Evaluation

Pre-Implementation

Implementation

- WITH clause performance issue
- Latch issue while gathering the stats
- Deadlock issue with HCC feature
- Informatica bulk load issue while loading from single source to multiple targets
- Wrong results with UNION ALL operator
- Query Issues

Questions?

Contact: ashish.mirji@saama.com , winston.pinto@saama.com



Thank You