

# Intel® Cloud Builders: Consumer Retail Website - Product Evaluation



This Blueprint is for a Consumer Retail Website application incorporating aspects of Business Intelligence (BI) based on a VMware vSphere\* Virtualization solution that leverages a Policy-Based Power Management Strategy to right-size the environment in correlation to its load. In addition, there are several design factors that predicate an understanding of the patterns of the Business Intelligence workload in question and how that workload behaves. The **most significant patterns** are called out for this application and are listed here by family, pattern name, a description, what problem the pattern solves (problem), key design decisions that influence the use of this pattern (driving forces), the typical participant patterns that this architectural pattern will use to solve the problem suggested by the scenario (collaborators), aspects of design than can be varied as a result of using this pattern (aspects that can vary), and the tradeoffs and results of using the pattern in terms of its limitations and constraints (tradeoff & constraints). This information is seen in the table below.

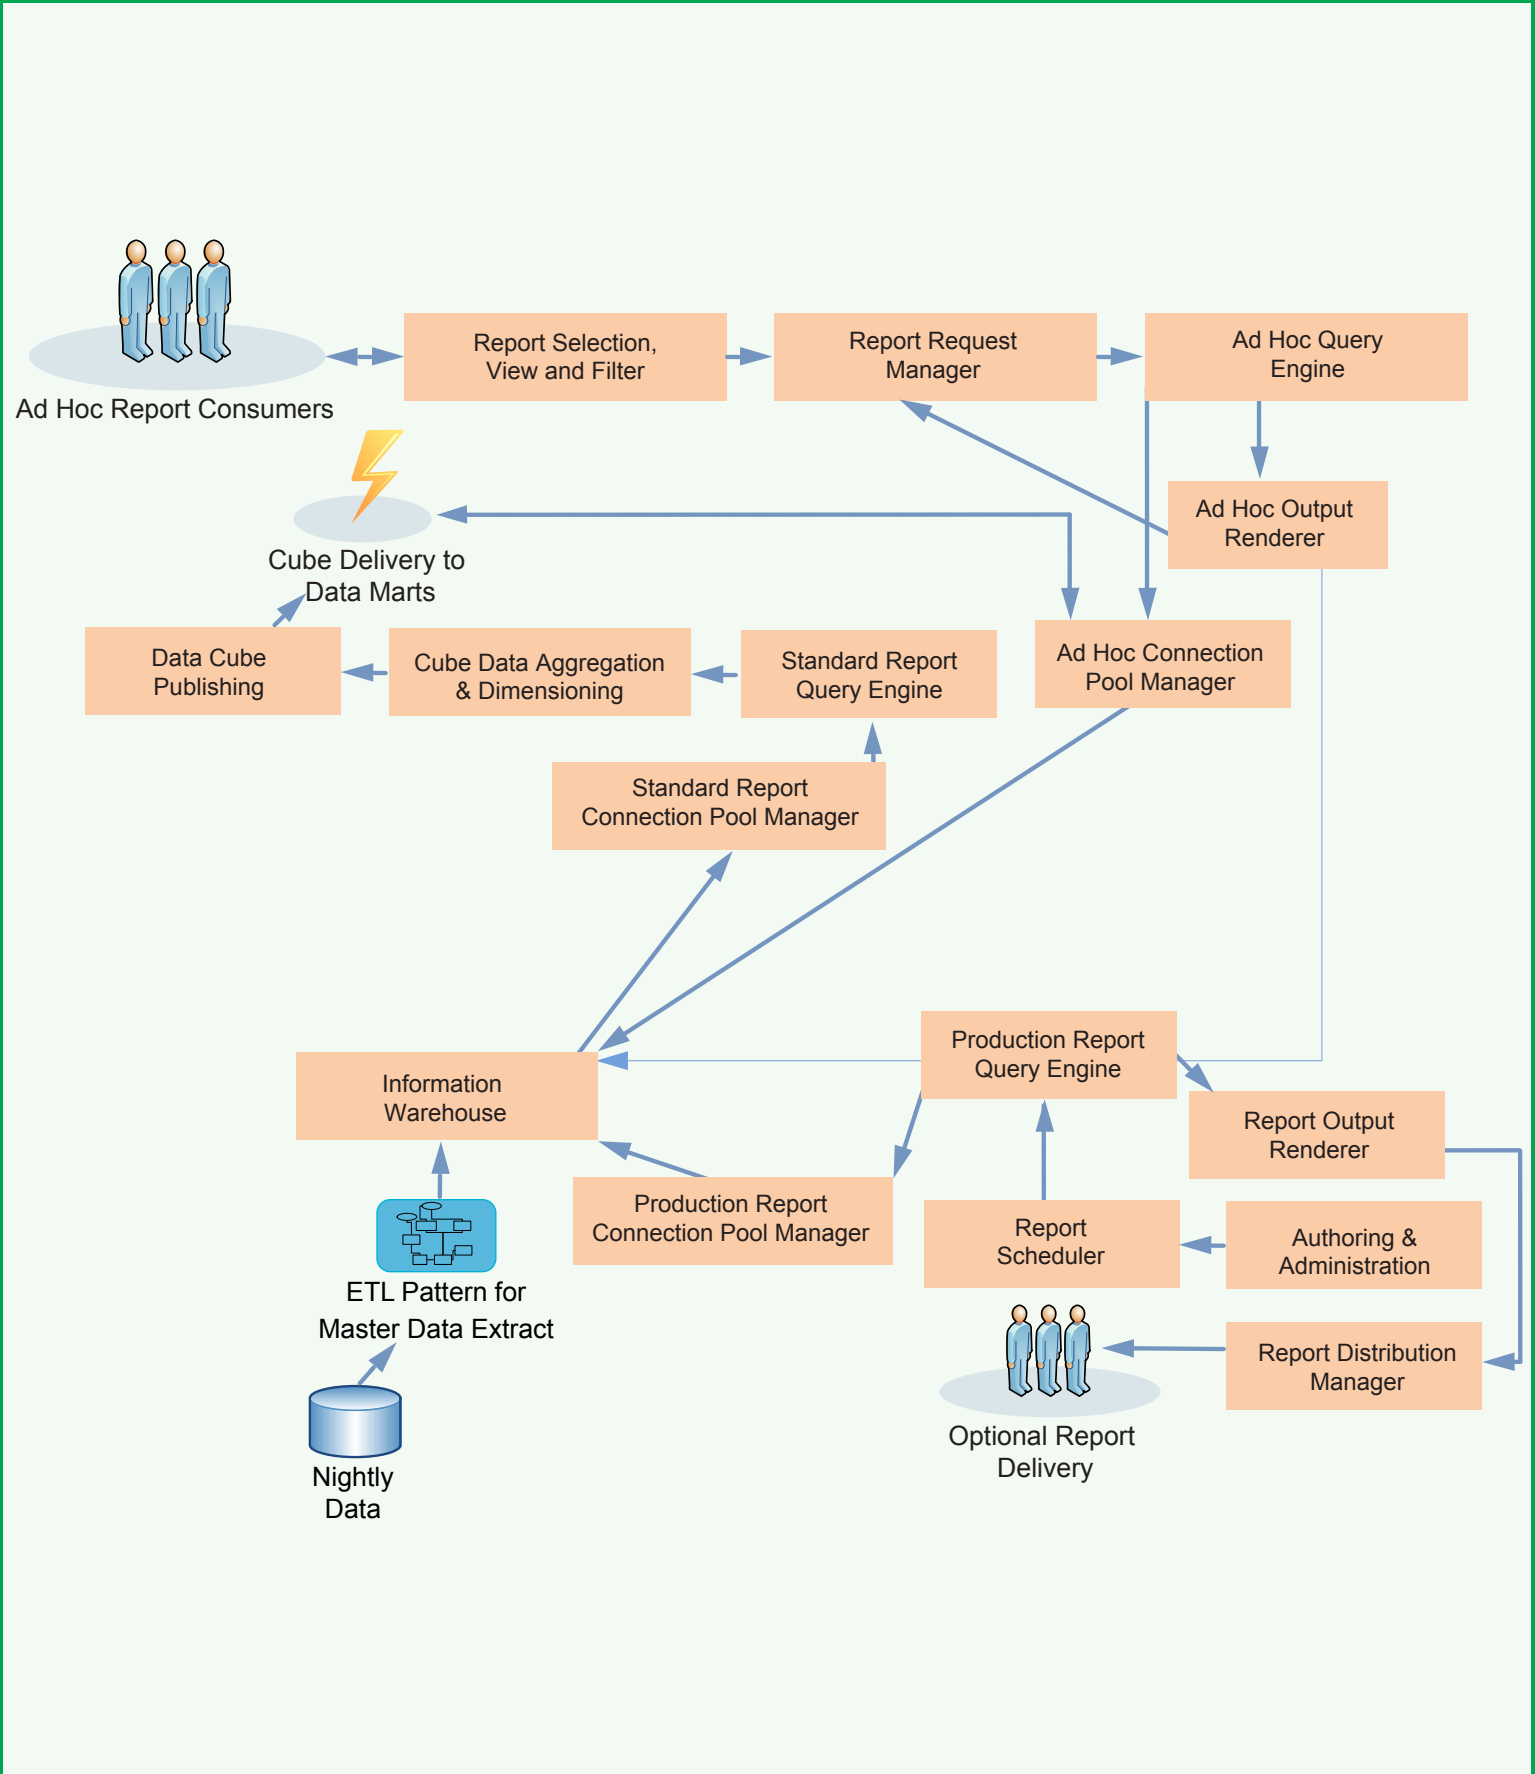
Using this knowledge, the following Blueprint sheets were generated by first considering the size of the workload to be applied: then performance requirements were used to generate virtual and logical views of the architectural, management, and physical infrastructure components needed to deploy this application in the cloud.

Instead of relying on the isolated intuition of architects and engineers to design the solution for cloud enablement, these blueprints are provided to ensure a more accurate and precise design is used as an initial instantiation to save on design, pilot and ultimately rebuild costs; and to enable more rapid go to market.

Pattern Family	Pattern Name	Brief Description	Problem	Driving Forces	Collaborators	Aspects that Can Vary	Tradeoffs & Constraints
Analytic System	Data Aggregator	Designed to aggregate many sources of data into pre-configured information hierarchies, categories or record types. This pattern will typically summarize existing information or collect data from many sources in order to transform or display it in a uniform matter. The performance of this app type pattern is characterized in the qualities (e.g. real time, batch) and not part of the canonical definition.	There is a need to aggregate many sources of data into pre-configured information hierarchies, categories or record types. The data might need to be transformed in order to summarize the disparate sources, making it available for display in a cohesive structure.	1) Multiple data sources have little in common with regard to structure and access mechanisms. 2) Multiple aggregation strategies are needed for different consumers. 3) Data qualities vary per input, and consumers have different data quality requirements. 4) Different consumers have unique delivery requirements.	A Data Aggregator pattern will be used when the aggregation problem is complex, and therefore separation of concerns is an important part of the design. Data Aggregators would call other patterns as a service in order to complete its tasks. Likely collaborators: a) Data Transformation, b) Data Driven Matcher (for reconciliations), c) Numerical Processor (for intensive calculations before summations), d) Portal Server - (for a comprehensive UI, when many sources and configuration options apply), e) Workflow pattern (for scheduling many complex aggregations), and f) Thick Client Portal would be client of a Data Aggregator.	1) Number of data sources. 2) Input formats. 3) Aggregation Structures. 4) Delivery service levels. 5) Data Aggregation Algorithms.	1) Multiple consumers and multiple sources, will increase the operational complexity, requiring scheduling or workflow. 2) Throughput will be a concern for aggregations with complex data structures and high volumes, solving these can increase operational complexity 3) Aggregations requiring very fast turnaround times may not be able to be mixed with long running aggregations and may require separate pattern instances. 4) Failover considerations get more complex for large data sets and/or complex hierarchies
Analytic System	Numerical Processor	Designed to optimize numerical calculations such as risk, pricing etc., this pattern specializes in processing numerical tasks such as multiple iterations of an algorithm. This pattern can perform calculations on large data sets with options for execution approach, Quality of Service levels and scenario choices. The Performance characteristics of Real-time/On Demand and batch are elicited in the qualities and are not part of the canonical definition.	There is a need to perform calculations on large data sets with options for parallel or serial execution; options for Quality of Service levels (e.g. response time, iteration level), and environment choices (to run scenarios under a variety of assumptions).	1) Multiple calculations will need to be performed simultaneously for different requestors. 2) Each calculation request will have a different data environment with its own directions for completion of the calculation. 3) Some calculations will have very high performance calculation requirements.	This pattern will collaborate with other patterns if data needs to be transformed prior to the calculations or aggregated or rendered after calculations. Possibly called by a) Data Aggregator (b) Thick Client Portal (c) Blackboard, (d) Event Driven Analysis & Response UI—may call (e) Transformation Engine.	1) Calculation iterations. 2) Service Level parameters that guide a when a calculation is good enough. 3) Scenarios. 4) Environments that scenarios run in.	1) Extreme Latency requirements will probably force the creation of a separate instance of a numerical processor. 2) If algorithms need to be parallelized then a grid solution will be required
Data Retention System	Data Warehouse	A Data Warehouse is a subject-oriented, integrated, time variant, nonvolatile collection of data in support of management decision-making process. OLAP provides one type of visualization mechanism, supporting multi-dimensional views because OLAP retains transformed data in a multi-dimensional cube for complex queries.	There is a need for a repository of consistent (not disparate) historical data that can be easily accessed and manipulated for decision support. This repository is needed to enable the understanding of patterns, trends and relationships in historical data by providing the foundation for enhanced visualization and decision support.	1) Highly indexed, heuristically tuned, derived historical data. 2) Structure that optimizes multi-dimensional queries. 3) Retain metadata 4) Visualization of consistent historical data.	Data Warehouses are not the owners of operational data—this pattern collaborates with analytic aggregation and transformation engines to obtain the data in the desired form.	1) Types of visualization. 2) Number of dimensions.	1) Query Performance is the biggest concern as the number of dimensions grows large.
User Interface	Portal Server	The nature of the response, and the degree of input vs. output in a Portal Server is project determined (e.g. Real Time, On Line, static, versus active transactional) and will be captured in the qualities, not the definition of the pattern.	There is a need for a presentation coordinator, acting on behalf of a set of clients that sends requests to and receives data from numerous service providers.	One client request will often decompose into multiple requests to disparate providers which may be self-contained systems that could require independent security validation. The PS must expect that the responses to this one request will return asynchronously in different formats. These formats will most likely have to be translated. The PS must be able to determine the minimal acceptable set of responses required before it is able to send a response to a client.	The Portal Server will collaborate with transformation, aggregation and formatting patterns in order to fulfill some requests.	1) Number of clients. 2) Number of providers. 3) Asynchronous processing of requests before a response is sent to the client. 4) Types of format translation.	The Portal Server must be have flexibility to accept different formats and providers, but still process requests in a timely manner. There will be major tradeoffs with throughput, response time, and flexibility of translation.
Reporting	OLAP & Ad Hoc Report Generator	OLAP augments the standard two-dimensional view of reports by allowing a user to compare rows within rows to columns within columns, effectively viewing multi-dimensional properties. Users can flip rows and columns; or invert the innermost columns/rows to the outermost and vice-versa. As a result of these capabilities, OLAP is considered ad hoc, because the users have a lot of discretion over how to view the data. OLAP reporting can generate data and conclusions without the benefit of direct users, by using program techniques. Data is stored in a different manner from an RDBMS. OLAP data is stored in a multi-dimensional cube, that often needs storage optimization	There is a need to allow users to create complex ad hoc multi-dimensional searches, one of which is typically time against an arbitrarily large data set. The need is to allow users to have a lot of discretion over how they view the data, switching rows within rows to columns within columns.	1) Well designed metadata. 2) Data cleansing process before the cube is built. 3) Speed of query for ad hoc. 4) Speed of canned report creation.	In order to meet the requirements for OLAP, a number of collaborations with other application patterns must occur, including the Application Integration Family of Patterns to extract and load data, and the Analytic Family of Patterns to translate, and transform.	1) Views along dimensions.	The major design concern is how to manage the potential for run away ad hoc queries.



# BI - Cloud Deployed Components



## BASELINE SELECTIONS

### Demand Characteristics

Total # of Records (#)	450M to 900M
Runtime SLA (hours)	Less than 4 hrs
Total Data Load (TB)	750GB to 1.5TB
Total Records Processed per Second (MB/s)	6MB/s to 12MB/s

### Resource Management Strategy

Trusted Compute Pools:	No
Policy-based Power Management	Yes

### Virtualization Management Strategy

VMware vSphere*:	Yes
Red Hat* Enterprise Virtualization:	No

## COMPUTE SELECTIONS

Processor Chipset: Intel® Xeon® processor X5570 Intel® Xeon® processor X7560 Intel® Xeon® processor L7555	Quantity: 1 6 1
Memory Modules: DDR3	Quantity: 57
Network Adapter: Intel® 82598EB 10 Gigabit Ethernet Controller	Quantity: 8
Disk Type: Disk Drive Module - SSD	Quantity: 27

## Pattern Function Descriptions

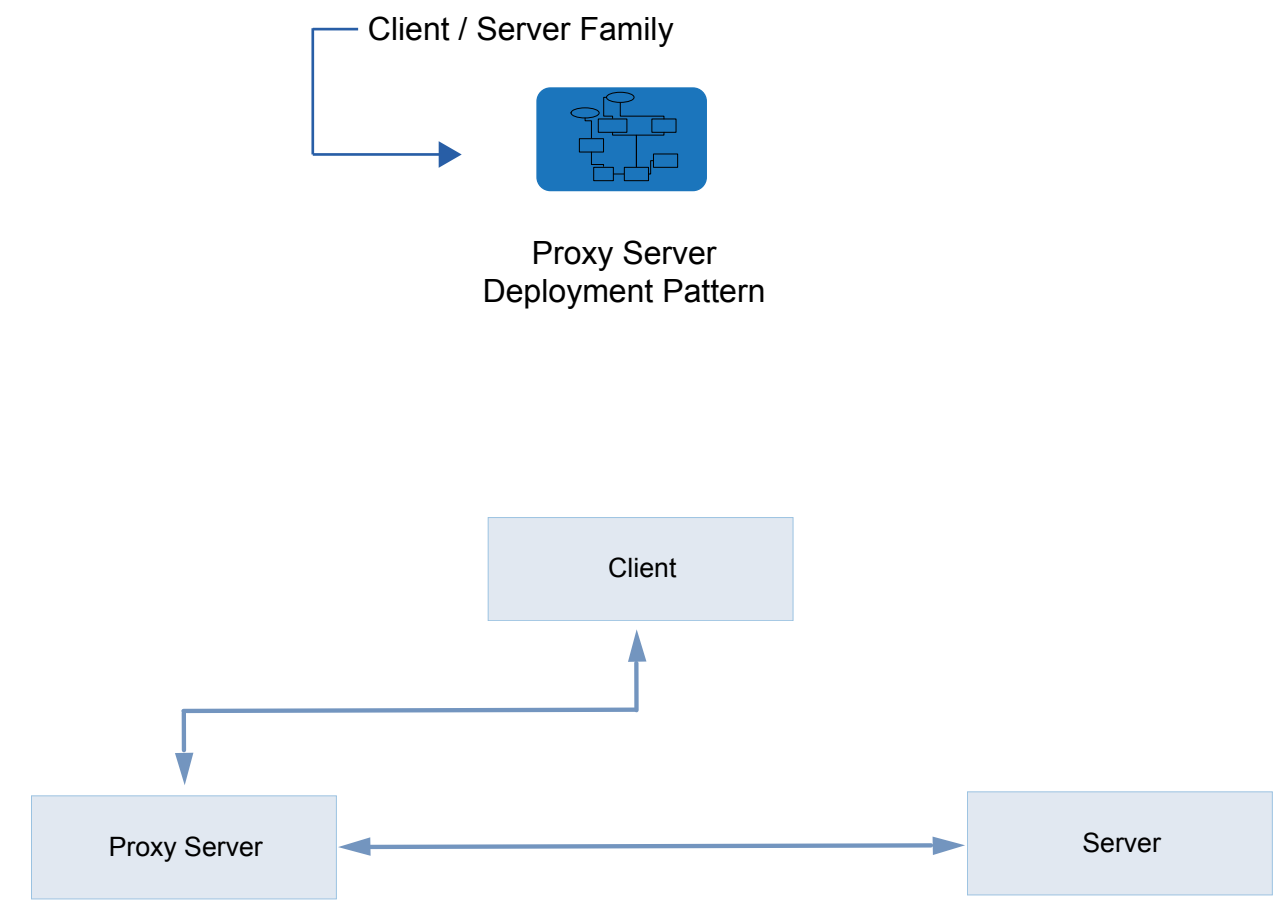
- Ad Hoc Report Output Renderer:** Creates customized views for ad-hoc queries
- Ad-Hoc Connection Pool Manager:** Provides connections to data marts for consumers in real time
- Ad-Hoc Query Engine:** Translates requests into actionable queries
- Thin Client Portal:** Channels Service Requests, holding partial responses till data is complete
- Data Mart:** Analytical data store designed to focus on specific business functions for a specific community within an organization.
- Cube Data Aggregator / Dimensioner:** Collects data and generates cube with specific dimensions
- Data Cube Publisher:** Pushes cube results to data marts
- Information Warehouse:** Holds views, stored procedures, and fact data
- Content Distributor:** Ensures consistent rules-based content delivery
- Production Report Connection Pool Manager:** Provides connections to the information warehouse for the production report generation
- Production Report Query Engine:** Translates requests into actionable queries
- Production Report Query Renderer:** Creates reports based on defined views
- Report Authoring & Administration:** Enables users to define query and report structure
- Report Distribution Manager:** Drives the delivery of reports
- Report Request Manager:** Processes requests for reports
- Report Scheduler:** Schedules delivery of reports based on lists and events
- Report Selector, Viewer, Filter:** Allows for report tailoring
- Standard Report Connection Pool Manager:** Allocates connections for report queries
- Standard Report Query Engine:** Translates requests into optimal queries

## Blueprint GPS

Shows a logical functional layout of a pattern or application. Also shows what the user selected for demand characteristics, compute, and storage



Illustrates a logical deployment architecture.



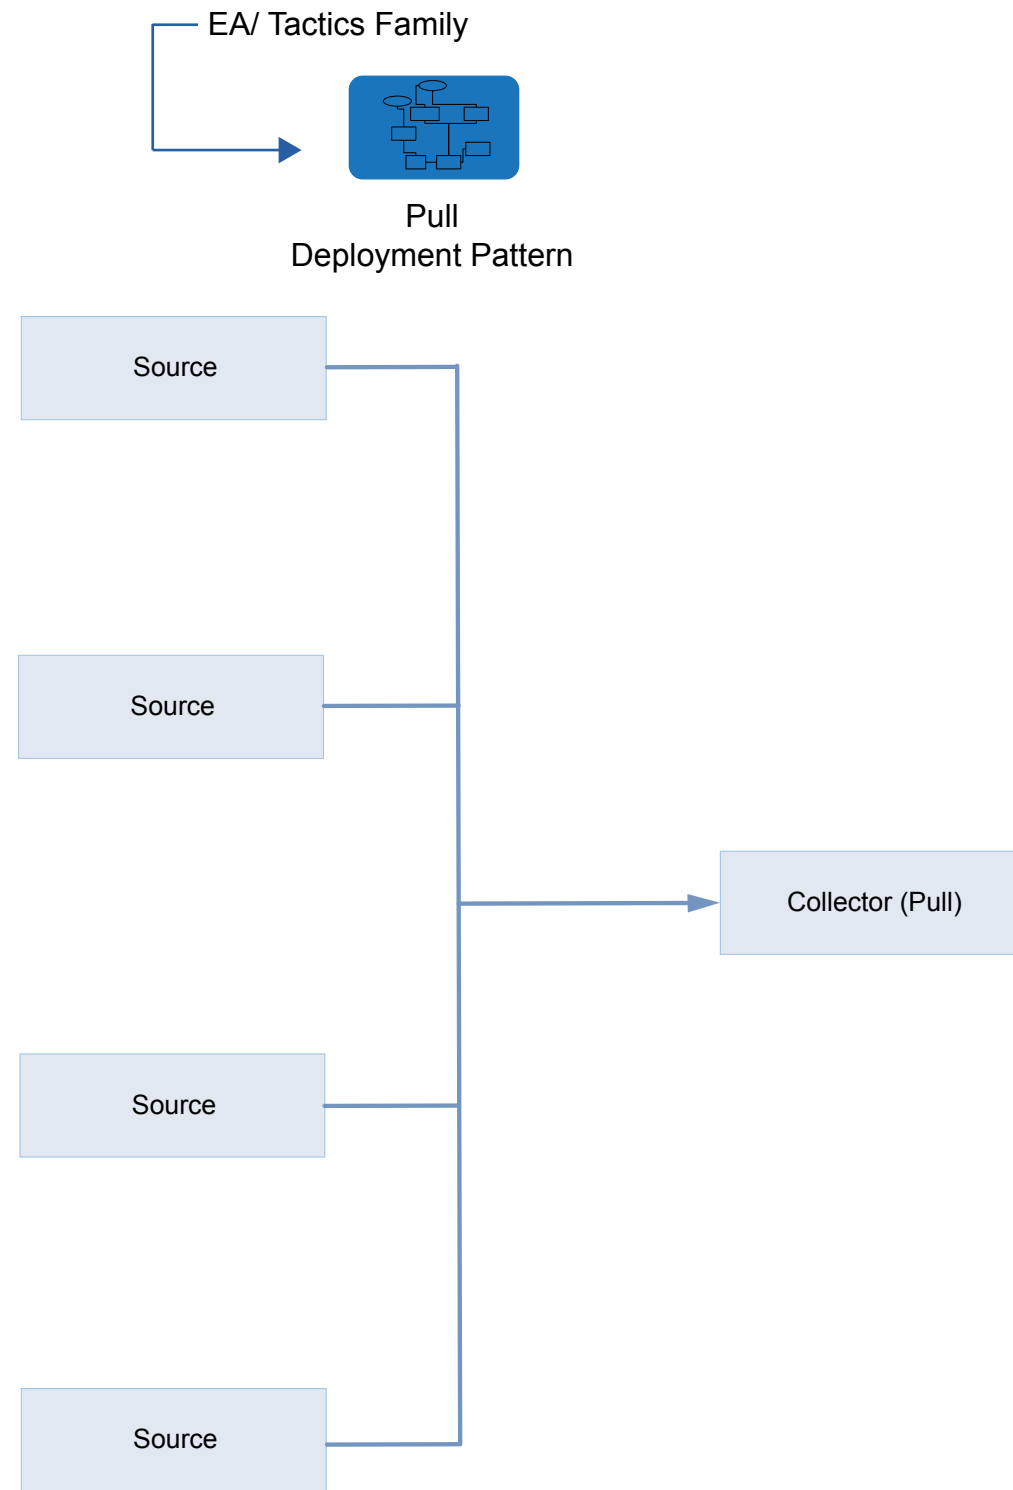
## Blueprint GPS

Shows which deployment pattern was used and the family of patterns that it came from. This is where a logical architecture would be deployed.

Note that more than one deployment pattern can be used to deploy a pattern or an application.

## Annotations

Illustrates a logical deployment architecture.



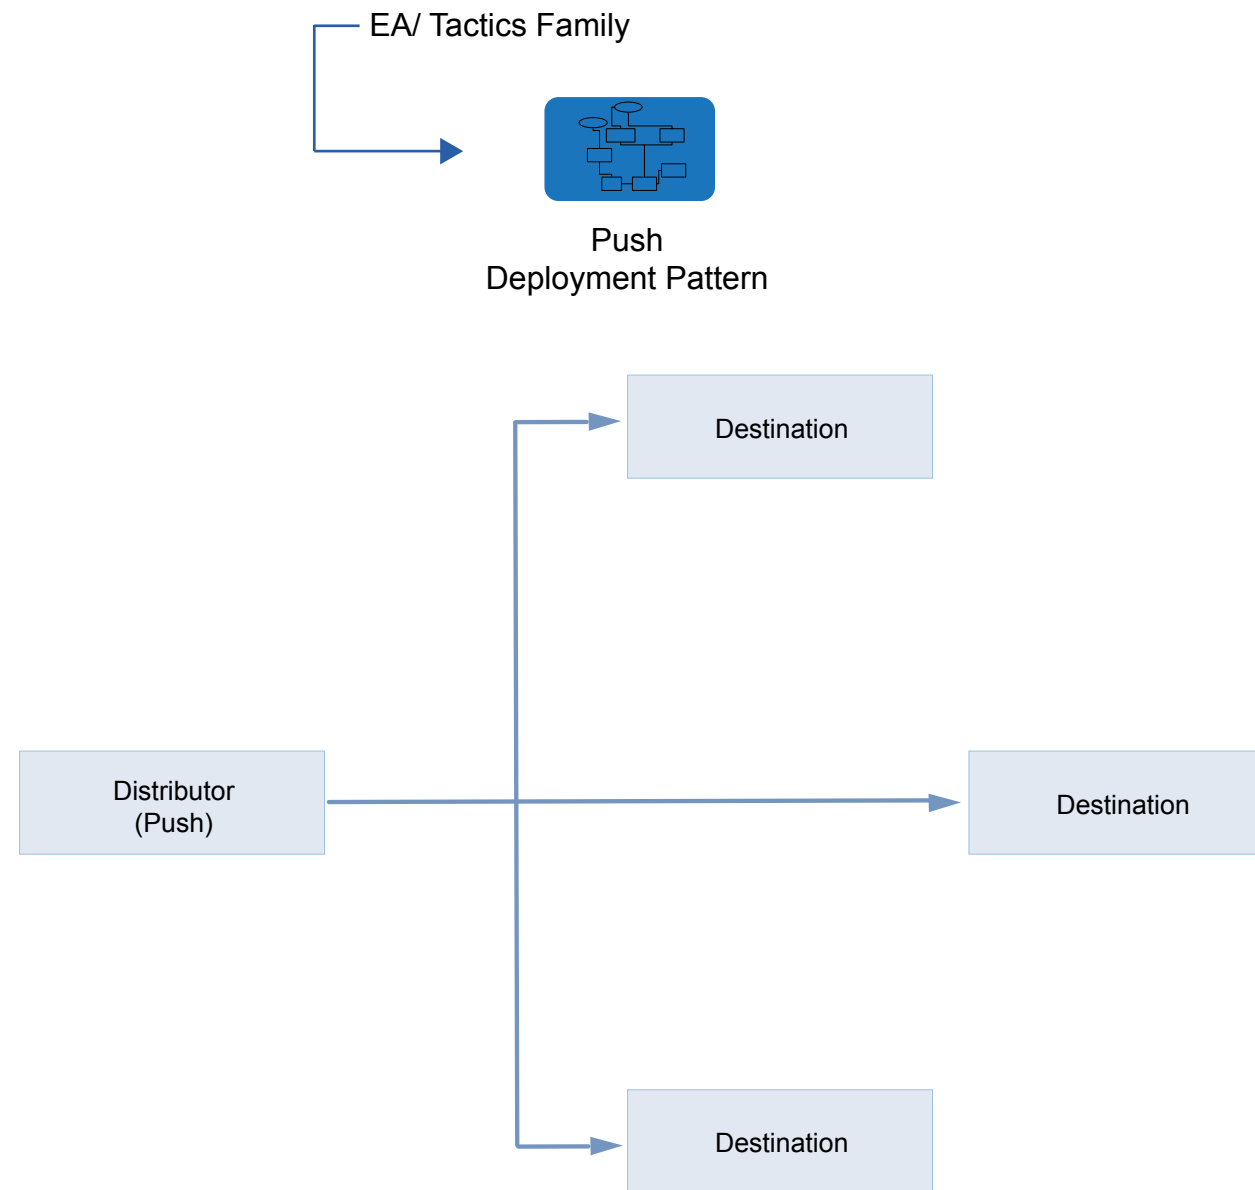
## Blueprint GPS

Shows which deployment pattern was used and the family of patterns that it came from. This is where a logical architecture would be deployed.

Note that more than one deployment pattern can be used to deploy a pattern or an application.

## Annotations

Illustrates a logical deployment architecture.



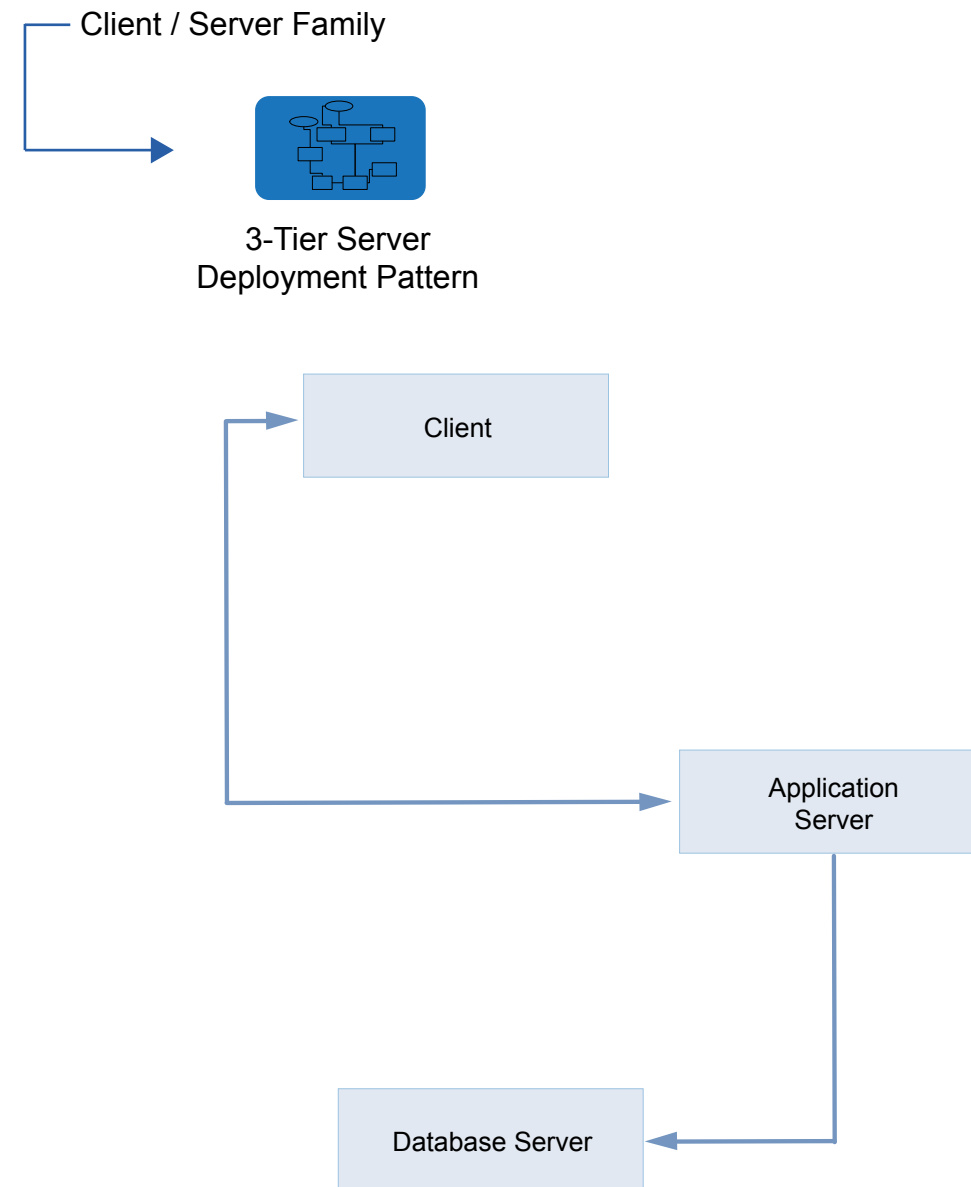
## Blueprint GPS

Shows which deployment pattern was used and the family of patterns that it came from. This is where a logical architecture would be deployed.

Note that more than one deployment pattern can be used to deploy a pattern or an application.

## Annotations

Illustrates a logical deployment architecture.



### Blueprint GPS

Shows which deployment pattern was used and the family of patterns that it came from. This is where a logical architecture would be deployed.

Note that more than one deployment pattern can be used to deploy a pattern or an application.

### Annotations



Pattern Component	Traditional Deployment Pattern Name	Deployment Pattern Component	Guest Virtual Machine Consumption Characteristics			
			Compute (GHz)	Memory (GB)	Disk (GB)	Network (Gbps)
Ad Hoc Report Output Renderer	3 Tier Server	Application Server	5.51	9.60	28.04	3.44
Ad-Hoc Connection Pool Manager	3 Tier Server	Application Server	2.36	17.28	98.13	3.44
Ad-Hoc Query Engine	3 Tier Server	Application Server	2.36	9.60	42.06	3.44
Cube Data Aggregator/ Dimensioner	Client Server	Server- Cube Creation	7.09	13.44	98.13	1.47
Data Cube Publisher	Push	Distributor Server	2.36	9.60	126.17	4.42
Information Warehouse	Pull	Collector Server	2.36	9.60	126.17	4.42
Production Report Connection Pool Manager	3 Tier Server	Database Server	5.51	17.28	126.17	4.42
Production Report Output Renderer	3 Tier Server	Application Server	5.51	17.28	126.17	1.47
Production Report Query Engine	Client Server	Server- renderer	7.09	17.28	126.17	4.42
Report Authoring & Administration	3 Tier Server	Application Server	1.58	3.84	28.04	0.98
Report Distribution Manager	Proxy Server	Proxy Server	2.36	9.60	126.17	4.42
Report Request Manager	Push	Distributor Server	2.36	9.60	126.17	4.42
Report Scheduler	3 Tier Server	Web Server	1.58	3.84	28.04	0.98
Report Selector, Viewer, Filter	Push	Distributor Server	2.36	9.60	42.06	3.44
Standard Report Connection Pool Manager	3 Tier Server	Web Server	5.51	17.28	126.17	4.42
Standard Report Query Engine	Client Server	Server- Cube Creation	7.09	17.28	126.17	4.42

## Configuration Notes:

The unit of work vectors, also called the consumption characteristics, provided above can be leveraged to construct the guest virtual machine instantiations necessary to deploy this application in the cloud. This organization of VMs by functional/application pattern component listed above is only one of numerous optimal deployments. In addition to this virtual layout, each VM will require additional configuration information. Additional configuration items for consideration are listed here:

1. <hostname> - This is the known DNS identifier and is widely published.
2. <ip\_address\_1> - This is the primary IP address used to locate or identify the system and this may be dynamic in nature.
3. <ip\_address\_2> - This is the secondary IP address used to locate or identify the system and this may be dynamic in nature.
4. <virtual\_ip\_address> - This is the static virtual IP address used to locate or identify the system. This value will seldom change (if ever).
5. <rack\_location\_name> - This is the current physical location of the VM (virtual machine) using a unique blade or rack naming convention.
6. <chassis\_location\_name> - This is the current physical location of the VM (virtual machine) using a unique chassis naming convention.
7. <facility\_location\_name> - This is the current physical location of the VM (virtual machine) using a unique data center naming convention.
8. <VM\_server\_hostname> - The system image is managed by a host server and this host server has a unique name associated with it. This location is also the CURRENT location of the boot kernel for the system image.
9. <guest\_os\_vendor> - This is the vendor OS type.
10. <guest\_os\_version> - This represents the version of the installed OS type for this system image.
11. <server\_function> - This identifies the INTENDED use of the system and includes Production, Development, Test, Staging or DR.
12. <service\_profile\_name> - This is the name of the service profile and should be unique for this system or unique to a pool of similar systems.

# Intel® Cloud Builders: Consumer Retail Website: Logical Configuration Page 1



Capability Name	Vendor Name	Product Name	Model/Version	Quantity (#)	Base Attribute 1	Value	Unit	Base Attribute 2	Value	Unit	Base Attribute 3	Value	Unit
Compute Processor	Intel	Intel® Xeon®	processor X5570	1	Max Core Clock Speed	2.66	GHz	Max # of Cores	4	#	Max # of Threads	8	#
Compute Processor	Intel	Intel® Xeon®	processor X7560	6	Max Core Clock Speed	2.93	GHz	Max # of Cores	4	#	Max # of Threads	8	#
Compute Processor	Intel	Intel® Xeon®	processor L7555	1	Max Core Clock Speed	1.86	GHz	Max # of Cores	2	#	Max # of Threads	2	#
Compute Rack Server	Intel	Non-specific	S5500WB	8	Max # Processor Sockets	2	#	Max # of Memory Slots	8	#	Max Memory Capacity	128	GB
Compute Memory	Non-specific	Non-specific	DDR3	57	Max Memory Capacity	4000	MB	Memory Type	DDR3	Categorical	Max # of Memory Ranks	2	#
Compute Chassis	Intel	Rack Chassis	SC5650BRP	2	Max Backplane Throughput	9600	Gbps	Max # Blade Slots	6	#	Max # Rack Slots	0	#
Converged Network Adapter	Intel	Eth Server Adapter	Intel® 82576EB	0	Max Throughput	1	Gbps	Max I/O	8000	IOPS	Max # of Ports	4	#
Converged Network Adapter	Intel	Eth Server Adapter	Intel® 82598EB	8	Max Throughput	10	Gbps	Max I/O	40000	IOPS	Max # of Ports	2	#
Storage Subsystem	Non-specific	Non-specific	6 Tray-14 DDMs/Tray	2	Max Raw Capacity	84000	GB	Max Disk Drives	84	#	Max Throughput	276	Gbps
Disk Drive Module	Non-specific	Non-specific	DD Module - FC	0	Max Raw Capacity	300	GB	Max Rotational Speed	15000	RPM	Form Factor	3.5	in
Disk Drive Module	Non-specific	Non-specific	DD Module - SATA	0	Max Raw Capacity	500	GB	Max Rotational Speed	7200	RPM	Form Factor	3.5	in
Disk Drive Module	Non-specific	Non-specific	DD Module - SSD	27	Max Raw Capacity	64	GB	Max Rotational Speed	N/A	RPM	Form Factor	2.5	in
Disk Configuration	Non-specific	Non-specific	RAID	1	Max Utilization	70	%	Min # of Disks	4	#	RAID Type	RAID 6	Categorical
Patch Level Monitoring	Non-specific	Non-specific SW	v0.0	1	Agent-less Inventory	Yes	Categorical	Secure Communications	Yes	Categorical	Virtual Environment Sup.	Yes	Categorical
Hardware Monitoring	Intel	RMM	v3.0	1	Max # Systems Monitored	1	#	Sample Rate	1	#/s	Max # of System Probes	10	#
Thermal Improvement	Intel	Node Manager	v1.5	1	Thermal Threshold Value	Yes	Categorical	Thermal Budget	Yes	Categorical	Thermal Time Limit Value	Yes	Categorical
Infrastructure Monitoring	Non-specific	Non-specific SW	v0.0	1	Max # Systems Monitored	1024	#	Real-time Monitoring	Yes	Categorical	Availability Monitoring	Yes	Categorical
Golden Image Generation	VMware*	vSphere*	v4.1	1	Max # Images Supported	Unlimited	#	Custom Image Templates	Yes	Categorical	Custom Media Repository	Yes	Categorical
Policy Based Provisioning	VMware*	vCenter	v4.1	1	Aggregate Physical Res.	Yes	Categorical	Affinity Rules	Yes	Categorical	Max VM Server Support	32	#
Power Monitoring	Intel	Node Manager	v1.5	1	Power Threshold Value	Yes	Categorical	Power Budget	Yes	Categorical	Power Time Limit Value	Yes	Categorical
Power QoS Policies	VMware*	vSphere*	v4.1	1	Monitor Power Usage	Yes	Categorical	Server Level Power Control	Yes	Categorical	Policy-based Power Mgmt	Yes	Categorical
Heat Dissipation Monitoring	Intel	Node Manager	v1.5	1	Thermal Threshold Value	Yes	Categorical	Thermal Budget	Yes	Categorical	Thermal Time Limit Value	Yes	Categorical
Database Cluster Management	VMware*	vSphere*	v4.1	1	Monitor Cluster State	Yes	Categorical	Synchronize Databases	No	Categorical	Create DBSnapshots	No	Categorical
VM Patch Management	VMware*	vSphere*	v4.1	1	OS Support	Windows	Categorical	VM Server Patch Mgmt	Yes	Categorical	VM Guest Patch Mgmt	Yes	Categorical
Dynamic Resource Pools	VMware*	vSphere*	v4.1	1	Dynamic Res.Balancing	Yes	Categorical	Resource Monitoring Sup	Yes	Categorical	Power Mgmt Support	Yes	Categorical
Thin Provisioning	VMware*	vSphere*	v4.1	1	Compute Thin Provisioning	Yes	Categorical	Memory Thin Provisioning	Yes	Categorical	Network Thin Provisioning	No	Categorical
Virtual Network Monitoring	VMware*	vCenter	v4.1	1	Agent-less Monitoring	Yes	Categorical	Network Performance	Yes	Categorical	Dependency Mapping	No	Categorical
Virtual Storage Configuration Management	VMware*	vSphere*	v4.1	1	VM Guest Datastore Mgmt	Yes	Categorical	Storage Subsystem Control	No	Categorical	LUN Provisioning	Yes	Categorical
Virtual Disk Management	VMware*	vSphere*	v4.1	1	VM Guest Support	Windows	Categorical	VM Server Support	Local	Categorical	(null)	(null)	(null)
Virtual Machine Monitoring	VMware*	vCenter	v4.1	1	Agent-less Monitoring	Yes	Categorical	VM State Monitoring	Yes	Categorical	(null)	(null)	(null)
Virtual Networks	VMware*	vSphere*	v4.1	1	Max # Virtual Switches	16	#	Max # Virtual Ports	256	#	VLAN Tagging Support	Yes	Categorical
Virtual Machines	VMware*	vSphere*	v4.1	1	Max # VMs per Server	512	#	Max Memory per VM	256	GB	Max CPU per VM	8	#
Virtual Machine Snapshots	VMware*	vSphere*	v4.1	1	Max # Snapshots	Unlimited	#	Max # Con.Guest Builds	1	#	(null)	(null)	(null)
Virtual Resource Monitoring	VMware*	vCenter	v4.1	1	Agent-less Monitoring	Yes	Categorical	VM Compute Monitoring	Yes	Categorical	VM Storage Monitoring	Yes	Categorical
Cloud Self-Service Portal	VMware*	vCloud Director*	v1.0	1	Max Virtual AppSupport	1028	#	Virtual Media Support	Yes	Categorical	Virtual Catalog Support	Yes	Categorical
Cluster/Pool Balancing	VMware*	vSphere*	v4.1	1	Dynamic Res.Balancing	Yes	Categorical	Resource Monitoring Sup	Yes	Categorical	Power Mgmt Support	Yes	Categorical
Workload Orchestration & Management	VMware*	vCloud Director*	v1.0	1	Policy-based resource Mgmt	Yes	Categorical	Integrate Autom.Processes	Yes	Categorical	Exception Handling	Yes	Categorical

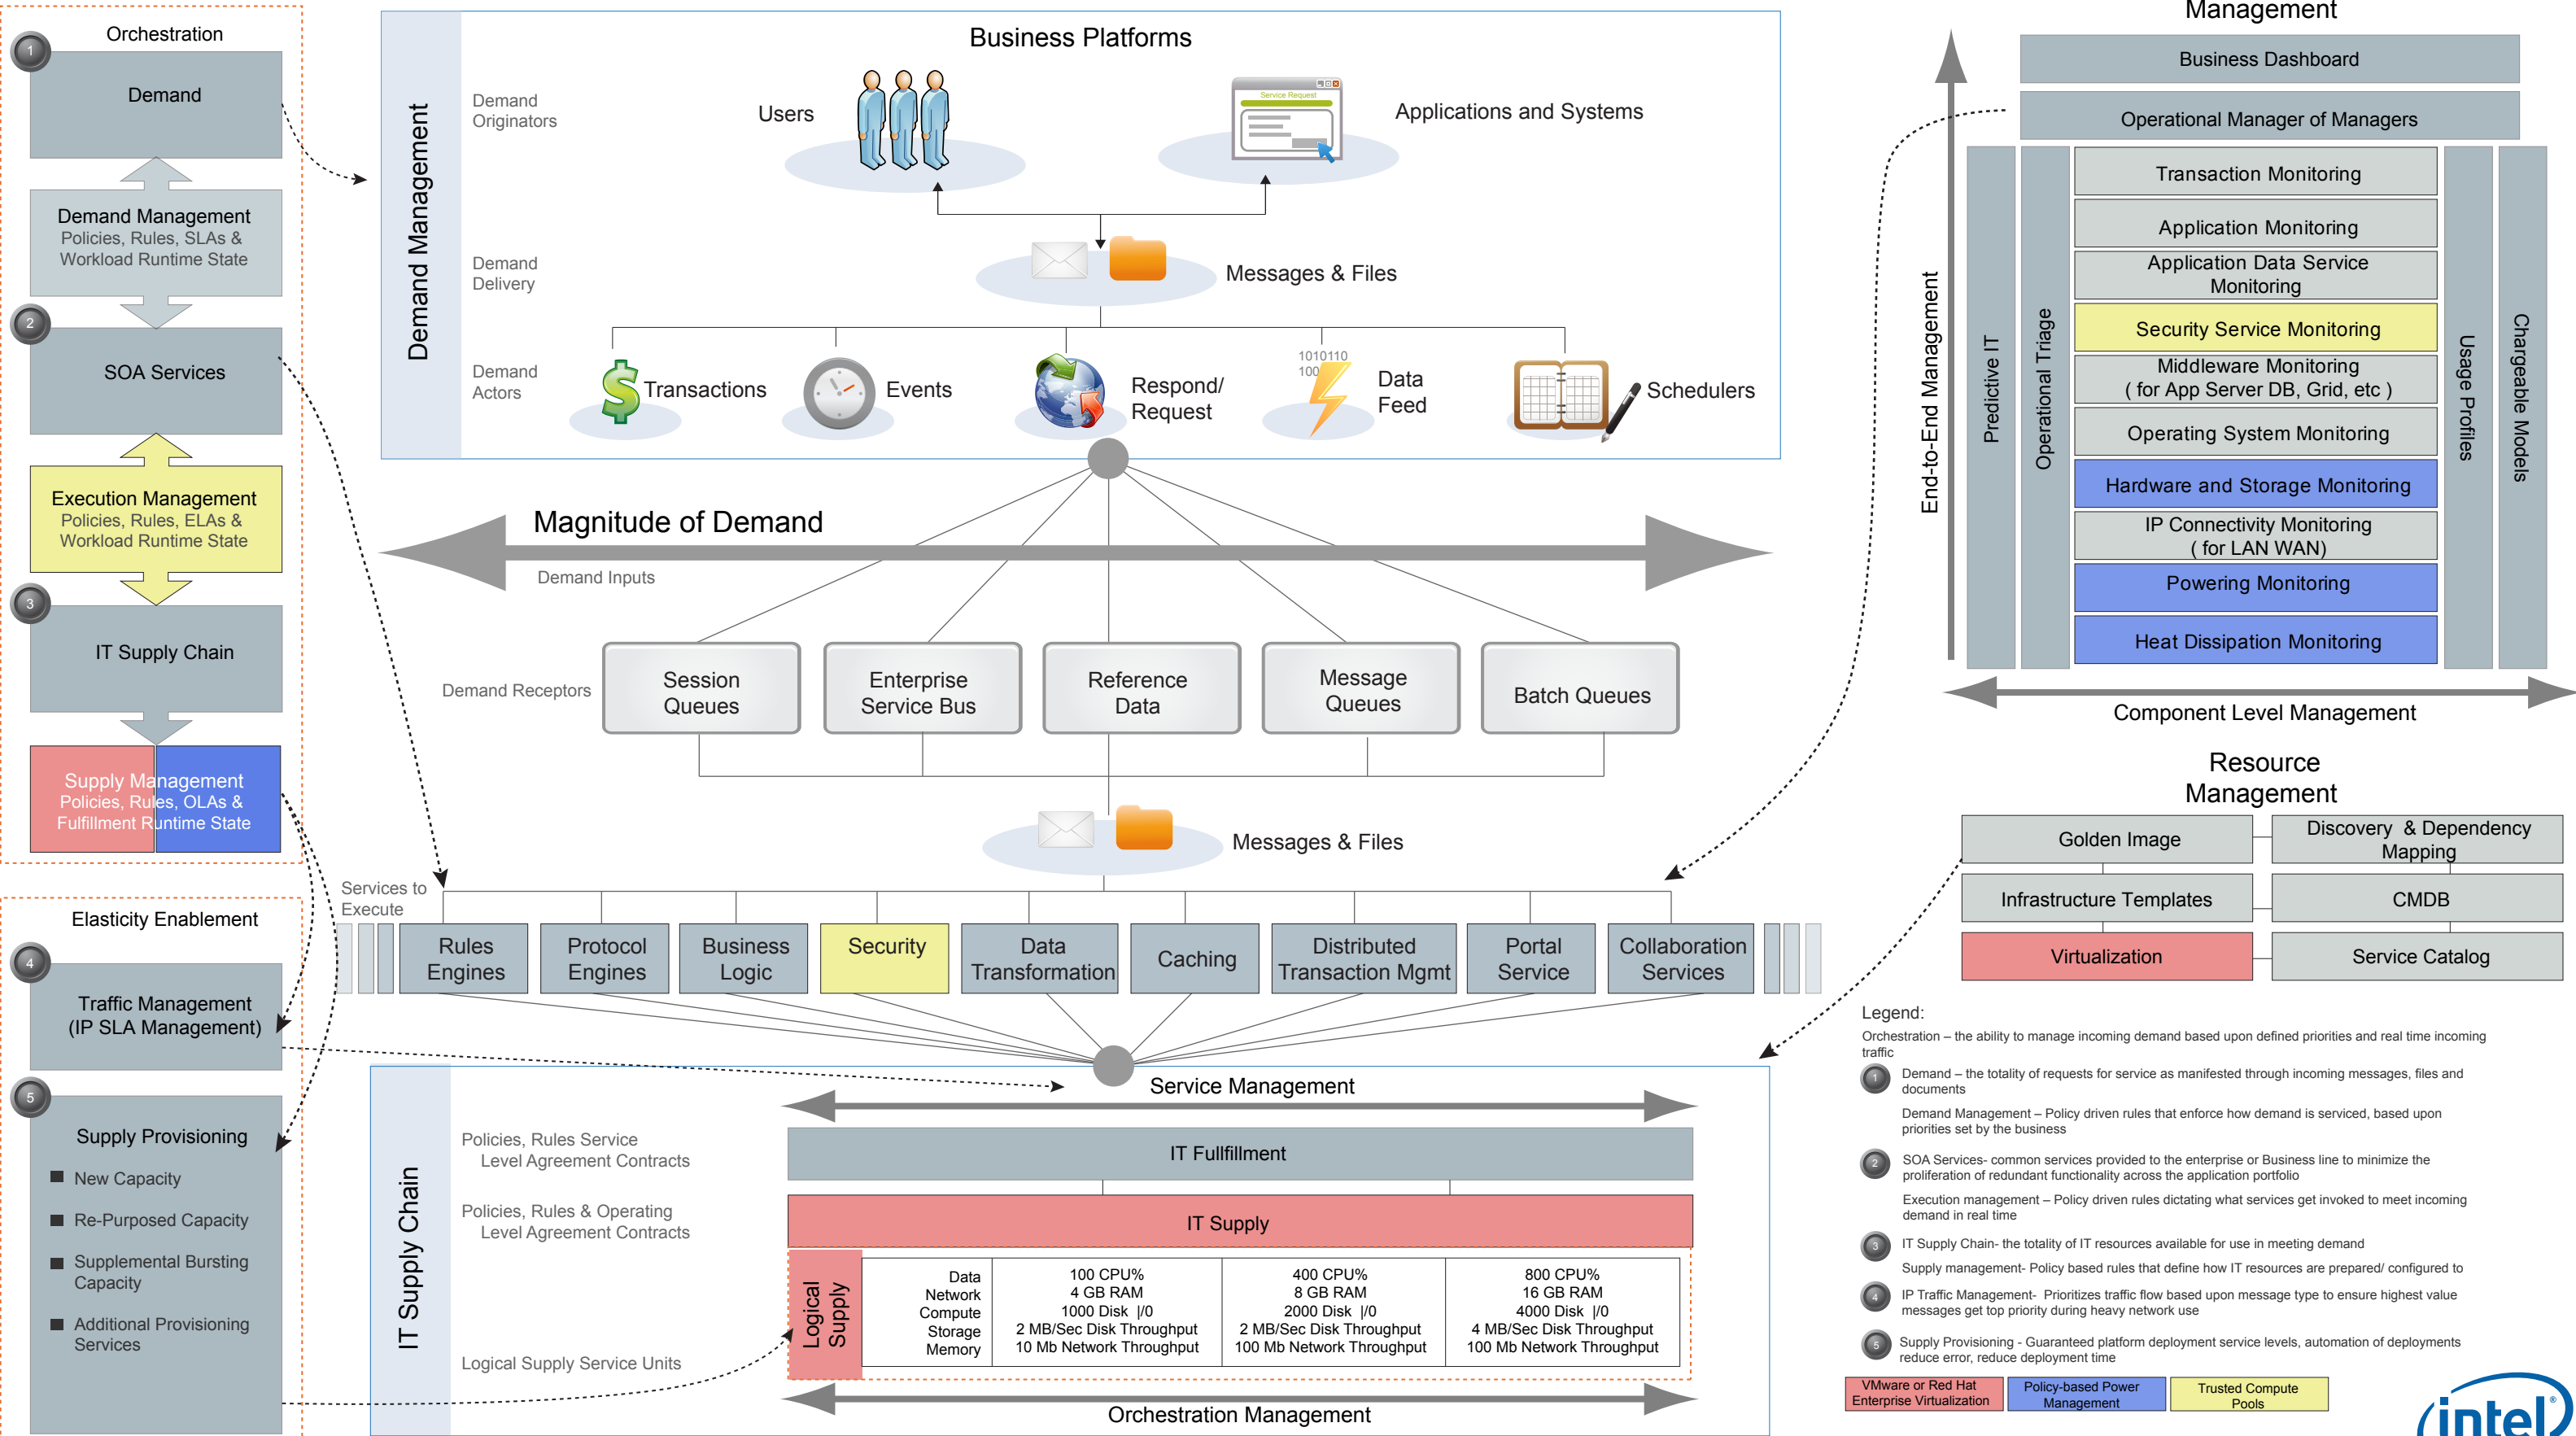




Capability Name	Vendor Name	Product Name	Model/Version	Quantity (#)	Base Attribute 4	Value	Unit	Base Attribute 5	Value	Unit
Compute Processor	Intel	Intel® Xeon®	processor X5570	1	Max Memory Size	144	GB	Bit Support	464	bits
Compute Processor	Intel	Intel® Xeon®	processor X7560	6	Max Memory Size	144	GB	Bit Support	64	bits
Compute Processor	Intel	Intel® Xeon®	processor L7555	1	Max Memory Size	144	GB	Bit Support	64	bits
Compute Rack Server	Intel	Non-specific	S5500WB	8	Max # of I/O Slots	6	#	Max I/O Bandwidth	40	Gbps
Compute Memory	Non-specific	Non-specific	DDR3	57	(null)	(null)	(null)	(null)	(null)	(null)
Compute Chassis	Intel	Rack Chassis	SC5650BRP	2	Max Power Consumption	1200	W	Height	6	U
Converged Network Adapter	Intel	Eth Server Adapter	Intel® 82576EB	0	(null)	(null)	(null)	(null)	(null)	(null)
Converged Network Adapter	Intel	Eth Server Adapter	Intel® 82598EB	8	(null)	(null)	(null)	(null)	(null)	(null)
Storage Subsystem	Non-specific	Non-specific	6 Tray-14 DDMs/Tray	2	Max Power Consumption	30380	W	Max Heat Output	96020	BTU
Disk Drive Module	Non-specific	Non-specific	DD Module - FC	0	(null)	(null)	(null)	(null)	(null)	(null)
Disk Drive Module	Non-specific	Non-specific	DD Module - SATA	0	(null)	(null)	(null)	(null)	(null)	(null)
Disk Drive Module	Non-specific	Non-specific	DD Module - SSD	27	(null)	(null)	(null)	(null)	(null)	(null)
Disk Configuration	Non-specific	Non-specific	RAID	1	(null)	(null)	(null)	(null)	(null)	(null)
Patch Level Monitoring	Non-specific	Non-specific SW	v0.0	1	(null)	(null)	(null)	(null)	(null)	(null)
Hardware Monitoring	Intel	RMM	v3.0	1	Max # of Process Probes	Unlimited	#	Max # of Log File Probes	10	#
Thermal Improvement	Intel	Node Manager	v1.5	1	Real-time HW Reading	Yes	Categorical	(null)	(null)	(null)
Infrastructure Monitoring	Non-specific	Non-specific SW	v0.0	1	Rules-based Automation	Yes	Categorical	Hist.Trend-based Reporting	Yes	Categorical
Golden Image Generation	VMware*	vSphere*	v4.1	1	Media Transcription	Yes	Categorical	Automatic Patch Insertion	Yes	Categorical
Policy Based Provisioning	VMware*	vCenter	v4.1	1	Max VM Support	1280	#	Policy-based Management	Yes	Categorical
Power Monitoring	Intel	Node Manager	v1.5	1	Real-time HW Reading	Yes	Categorical	(null)	(null)	(null)
Power QoS Policies	VMware*	vSphere*	v4.1	1	VM Scaling	Yes	Categorical	(null)	(null)	(null)
Heat Dissipation Monitoring	Intel	Node Manager	v1.5	1	Real-time HW Reading	Yes	Categorical	(null)	(null)	(null)
Database Cluster Management	VMware*	vSphere*	v4.1	1	Log Shipping	Yes	Categorical	Create Database Mirror	Yes	Categorical
VM Patch Management	VMware*	vSphere*	v4.1	1	(null)	(null)	(null)	(null)	(null)	(null)
Dynamic Resource Pools	VMware*	vSphere*	v4.1	1	(null)	(null)	(null)	(null)	(null)	(null)
Thin Provisioning	VMware*	vSphere*	v4.1	1	Datastore Thin Provisioning	Yes	Categorical	(null)	(null)	(null)
Virtual Network Monitoring	VMware*	vCenter	v4.1	1	(null)	(null)	(null)	(null)	(null)	(null)
Virtual Storage Configuration Management	VMware*	vSphere*	v4.1	1	(null)	(null)	(null)	(null)	(null)	(null)
Virtual Disk Management	VMware*	vSphere*	v4.1	1	(null)	(null)	(null)	(null)	(null)	(null)
Virtual Machine Monitoring	VMware*	vCenter	v4.1	1	(null)	(null)	(null)	(null)	(null)	(null)
Virtual Networks	VMware*	vSphere*	v4.1	1	(null)	(null)	(null)	(null)	(null)	(null)
Virtual Machines	VMware*	vSphere*	v4.1	1	Max I/O per VM	(null)	IOPS	Max Throughput per VM	(null)	Gbps
Virtual Machine Snapshots	VMware*	vSphere*	v4.1	1	(null)	(null)	(null)	(null)	(null)	(null)
Virtual Resource Monitoring	VMware*	vCenter	v4.1	1	VM Memory Monitoring	Yes	Categorical	(null)	(null)	(null)
Cloud Self-Service Portal	VMware*	vCloud Director*	v1.0	1	Virtual Organization Support	Yes	Categorical	Partitioned Network Sup.	Yes	Categorical
Cluster/Pool Balancing	VMware*	vSphere*	v4.1	1	Automatic Pool Balancing	Yes	Categorical	Policy Based Pool Balance	Yes	Categorical
Workload Orchestration & Management	VMware*	vCloud Director*	v1.0	1	Centralized Resource Mgmt	Yes	Categorical	Auditing	Yes	Categorical

# Intel® Cloud Builders Demand Driven Execution Management

Provides an overview introduction to Execution Management. It shows the Scope of Dynamic Infrastructure Management Capabilities that must be adopted to achieve a real-time infrastructure. It is expected that the organization would adopt these in phases using a top-down process



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at [www.intel.com](http://www.intel.com).

Copyright © 2010 Adaptivity, Inc. All Rights Reserved.

Copyright © 2010 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

\*Other names and brands may be claimed as the property of others.