

# Intel® Cloud Builders: Local Real Estate Website



This is a blueprint for a Local Real Estate Website implemented as a Rich Internet Application (RIA), based on a VMware vSphere\* Virtualization solution that leverages Trusted Compute Pools to ensure security and 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 that 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.

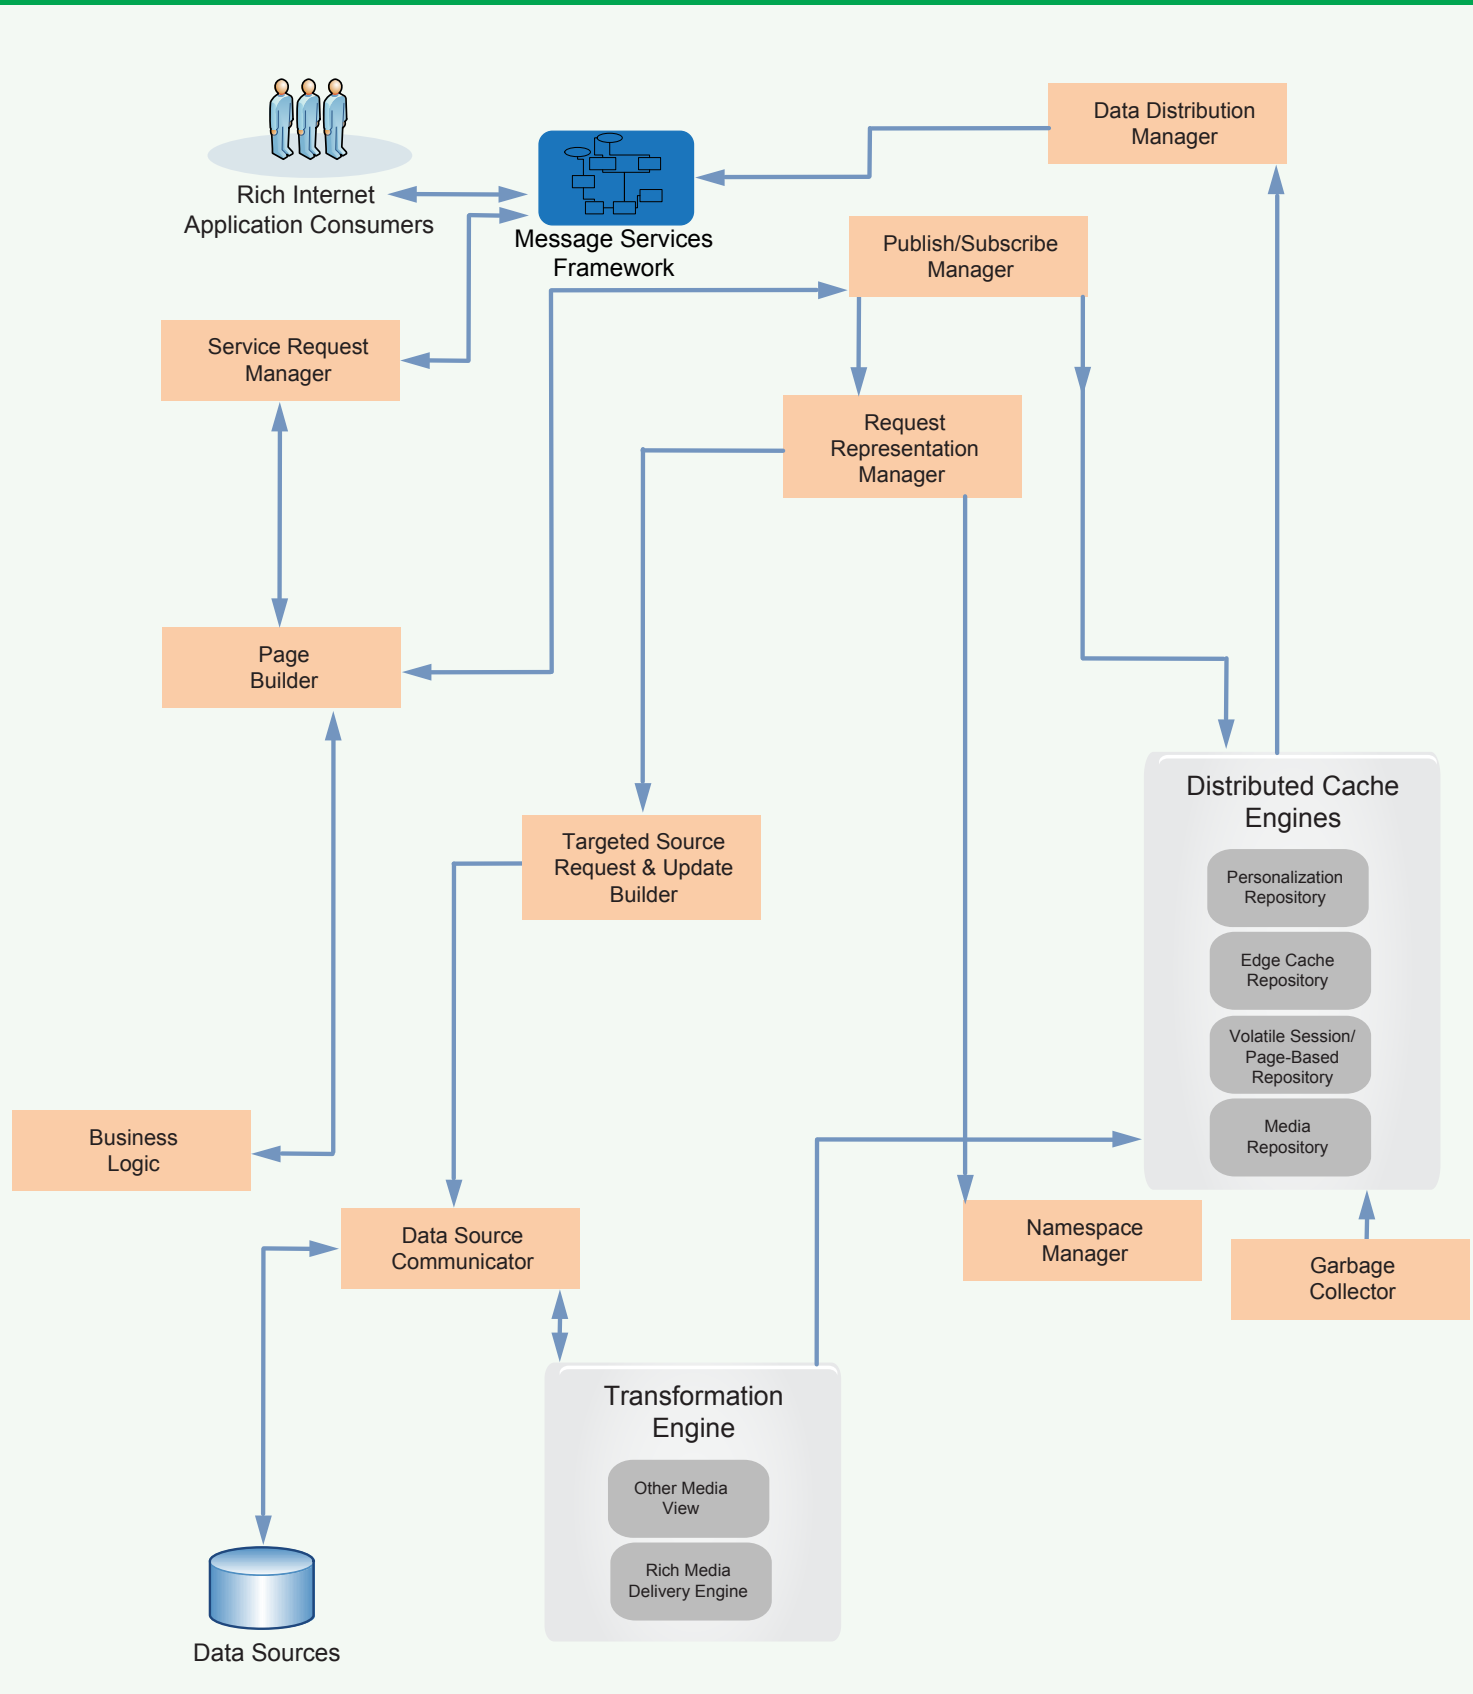
Using this knowledge, the following Blueprint sheets contained in this binder were generated by first considering the size of the workload to be applied and then the performance requirements was then used to generate a virtual and logical view of the architectural, management and physical infrastructure 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 already 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	Transformation Engine	Focus on the transformation of data between sets of representations. Input and output streams could be multiple in nature. This pattern typically utilizes reference data to retrieve transformation rules, and could employ context driven rules to execute transformations.	There is a need to transform multiple, arbitrarily sized and formatted data sets from one representation to another. The transformation rules can be varied, complex and are subject to frequent change.	1) There are multiple input and output streams. 2) The rules guiding the transformations can be context driven for each stream and tend to change frequently. 3) One input stream can be transformed to multiple output stream formats. 4) Each stream will have it's own service delivery option.	Transform engines will likely be called upon to perform a service by other patterns. Likely clients are a) Numerical Processor, b) Data Aggregator, c) Enterprise Service Bus, and d) Message Processor. Likely service collaborators are Data Driven Matcher	1) Number of data sources. 2) Input formats. 3) Output formats. 4) Delivery service levels. 5) Transformation Rules.	1) Multiple consumers and multiple sources, will increase the operational complexity, requiring scheduling or workflow. 2) Throughput will be a concern for complex transformations with high volumes, solving these can increase operational complexity. 3) Transformations requiring very low latency will not be able to be mixed with long running transformations and will require separate pattern instances. 4) Failover considerations get more complex for large data sets.
Data Retention System	Distributed Cache	Holds data in memory to expedite processing performance. Data held can be static, or highly dynamic. Aspects of a data set can be held across tiers, requiring synchronization / invalidation when data has changed. Data held can have representations that are independent of the persistent data model, and these representations can be independent from tier to tier.	There is a need to accelerate the throughput and/or response time for an application, by keeping frequently used objects in memory, so as to avoid expensive disk retrievals. These objects are subject to change from multiple sources, therefore these changes must be propagated to all users and all copies upon the change event.	1) Throughput needs to be maximized. 2) Local and distributed cache needs to be synchronized with ACID properties. 3) Information must be stored from different formats. 4) Cache tends to be highly volatile.	1) Other cache instances. 2) Transaction managers 3) Transactional databases.	Localized and Distributed Caching- Includes-cache consistency, synchronization, garbage collection, on demand loading, namespace access, object expiration & fault tolerance.	1) Caching does not guarantee performance enhancement it needs to be carefully designed. 2) The use of Caching increases the complexity of an application. 3) The use of caching increases the risk exposure for an application because failure recovery is more difficult.
Data Retention System	Transaction Data Base	Transaction-based data retention systems must have ACID properties in order to support operational processing of data records whose applications require high data integrity.	There is a need to have the results of a transaction be guaranteed, durable, non-refutable, and serve as the data of record for an application.	1) Integrity of the transactions must meet ACID properties. 2) Failure recovery must support the ACID principles. 3) Throughput will always be a consideration.	The typical patterns that would collaborate would be involved in the processing of transactions, such as workflow and transaction managers.	1) This would be vendor specific.	1) Reliability and throughput performance would need to be very highly rated over all other aspects.
User Interface	Portal Server	The nature of the response, and the degree of input vs. output in a PS 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 maybe 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.
Application Integration System	Enterprise Service Bus	An ESB is designed to abstract most of the mechanics of messaging from the application. The choice of protocols, formats, messaging behaviors will be configurable in an ESB. Applications will be able to specify the intent (e.g. pub/sub, guaranteed) of the interaction and there will be a pre-configured set of messaging features that the ESB can utilize to achieve the application goal. ESB represents the next evolution in integration among applications.	There is a need for applications to be able to specify the intent behind their message interaction goals (e.g. send to many recipients, guarantee the delivery) without having to code the specific protocol behaviors. An intermediary is required to align the service requirement with available mechanisms so that applications can decouple messaging logic and configuration tasks from business logic.	1) Performance. 2) Enterprise Critical High Availability. 3) Types of protocols that can be processed.	The ESB is a large-scale architecture pattern that will aggregate many of the other patterns in this Application Integration Family, including a variety of Brokers, Routers and Gateways.	1) Protocols. 2) Message delivery Service Levels. 3) Types of collaborators.	1) Performance and High Availability are absolutely critical for an ESB and will put constraints on the variety of protocols that the ESB can broker. 2) The ESB will most like serve in a widely dispersed geographical area, placing challenges on failover.
Transaction System	Distributed Transaction Manager	Transactions can be nested (multiple parts that can be processed independently or sequentially). The DTM marshals the transaction through its lifecycle of process steps, which could include processes on distributed machines. One of the key responsibilities of the DTM is to ensure ACID transactional properties in a distributed environment, which will include the ability to guarantee two-phase commit and rollback. It has a Transaction State Repository, State Transition Engine, Transaction Sequence Orchestration, Routing and an XA or Equivalent Resource Manager.	There is a need to coordinate complex transaction behavior between collaborating applications.	1) Reliability of transaction marshalling to ensure ACID properties. 2) Reconciliation strategies to resolve breaks due to failure.	It will collaborate with Event Processors, Data Repositories, and Message Brokers/Routers.	1) Degree of transaction nesting.	1) XA transactions have severe response time and throughput constraints. 2) XA transactions do not guarantee that transactions will meet all ACID properties in the case of failure and it is recommended that participants create reconciliation software to address anomalies.



## RIA - Cloud Deployed Components



## BASELINE SELECTIONS

### Demand Characteristics

Peak # of Concurrent Users	1,000 to 10,000
Peak Transaction Size (bytes/transaction)	5kB to 25kB
Peak Response Size (bytes/transaction)	1MB to 10MB
Peak Interval Between Transactions per User (s)	2s to 4s

### Resource Management Strategy

Trusted Compute Pools:	Yes
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 0
Memory Modules: DDR3	Quantity: 44
Network Adapter: Intel® 82598EB 10 Gigabit Ethernet Controller	Quantity: 8
Disk Type: Disk Drive Module - SSD	Quantity: 104

## Pattern Function Descriptions

**Service Request Manager:** Accepts requests; provides rich interaction facility through high speed delivery of large multimedia content; interprets partial page updates

**Business Logic:** Interacts with user; responds to partial page updates; invokes necessary logic for refresh

**Page Builder:** Waits for page elements; deliver partial page updates; delivery pages based on configuration; invoke business logic

**Garbage Collector:** Cleans expired entities

**Edge Cache Repository:** Holds more stable content that changes less and has a high reuse profile (e.g. Icons, banners, stock videos, form shells)

**Rich Media Transformation Engine:** Transform to required format; Deserialization / Serialization

**Non-Rich Media Transformation Engine:** Transform to required format; Deserialization / Serialization

**Targeted Source & Request Update Builder:** Build requests for each source; build updates for each source

**Pub/Sub Manager:** Process subscription requests; process publish requests; alert subscriber of change; bundles requests

**Data Distribution Manager:** Uses messaging semantics to guarantee delivery of changed data; has ACID transaction semantics

**Namespace Manager:** Ensures data entities are properly named, preserving uniqueness

**Request Representation Mapper:** Holds source metadata; resolves complex requests to source

**Data Source Communicator:** Makes requests and updates; awaits responses; awaits unsolicited updates from remote data changes; has API's to refine types of requests

**Session Page Based Repository:** Holds page element based data for streamlined session throughput; higher volatility; holds intermediate values during sessions

**Media Repository:** Holds media that can be streamed to multiple clients

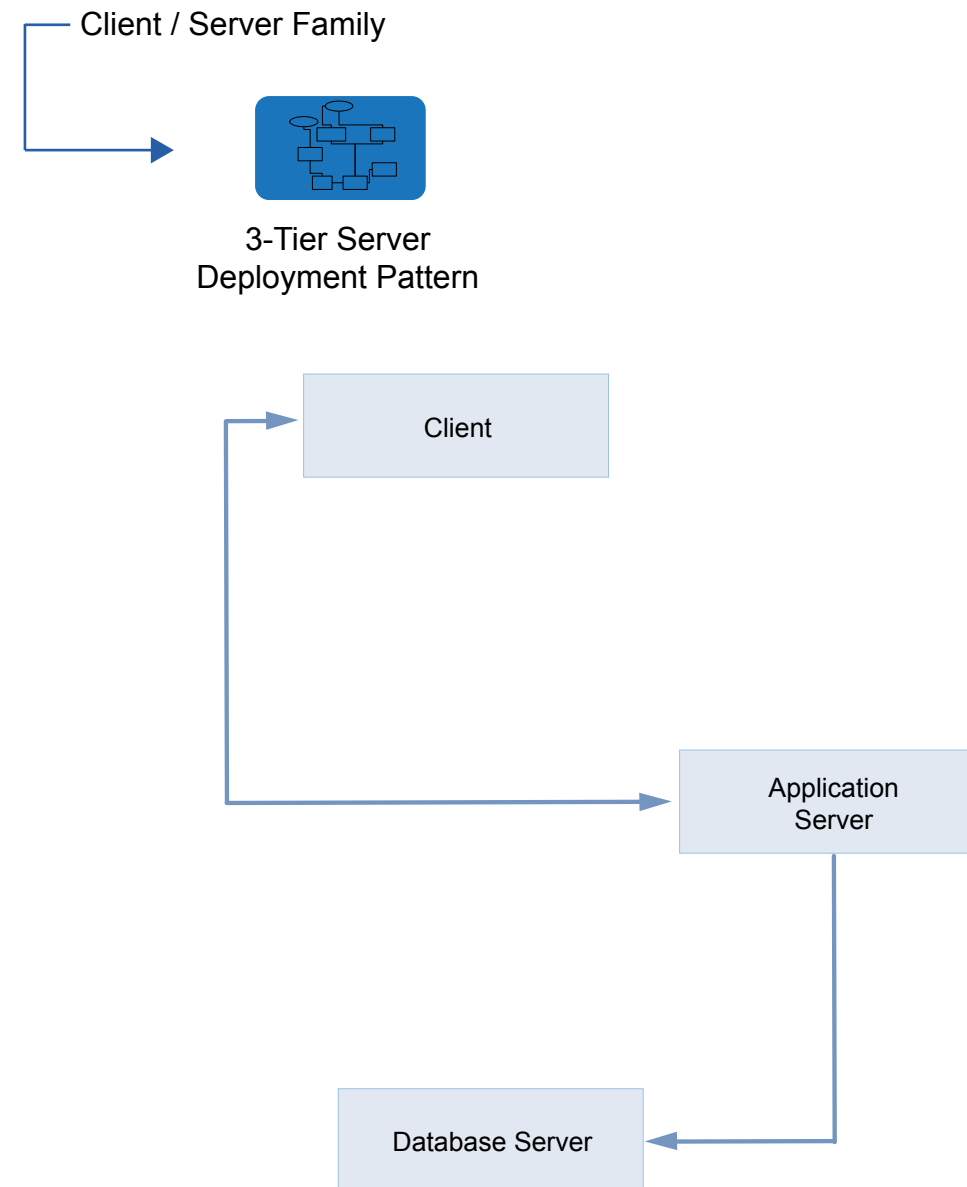
**Personalization Repository:** Holds views; holds format preferences

## 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.

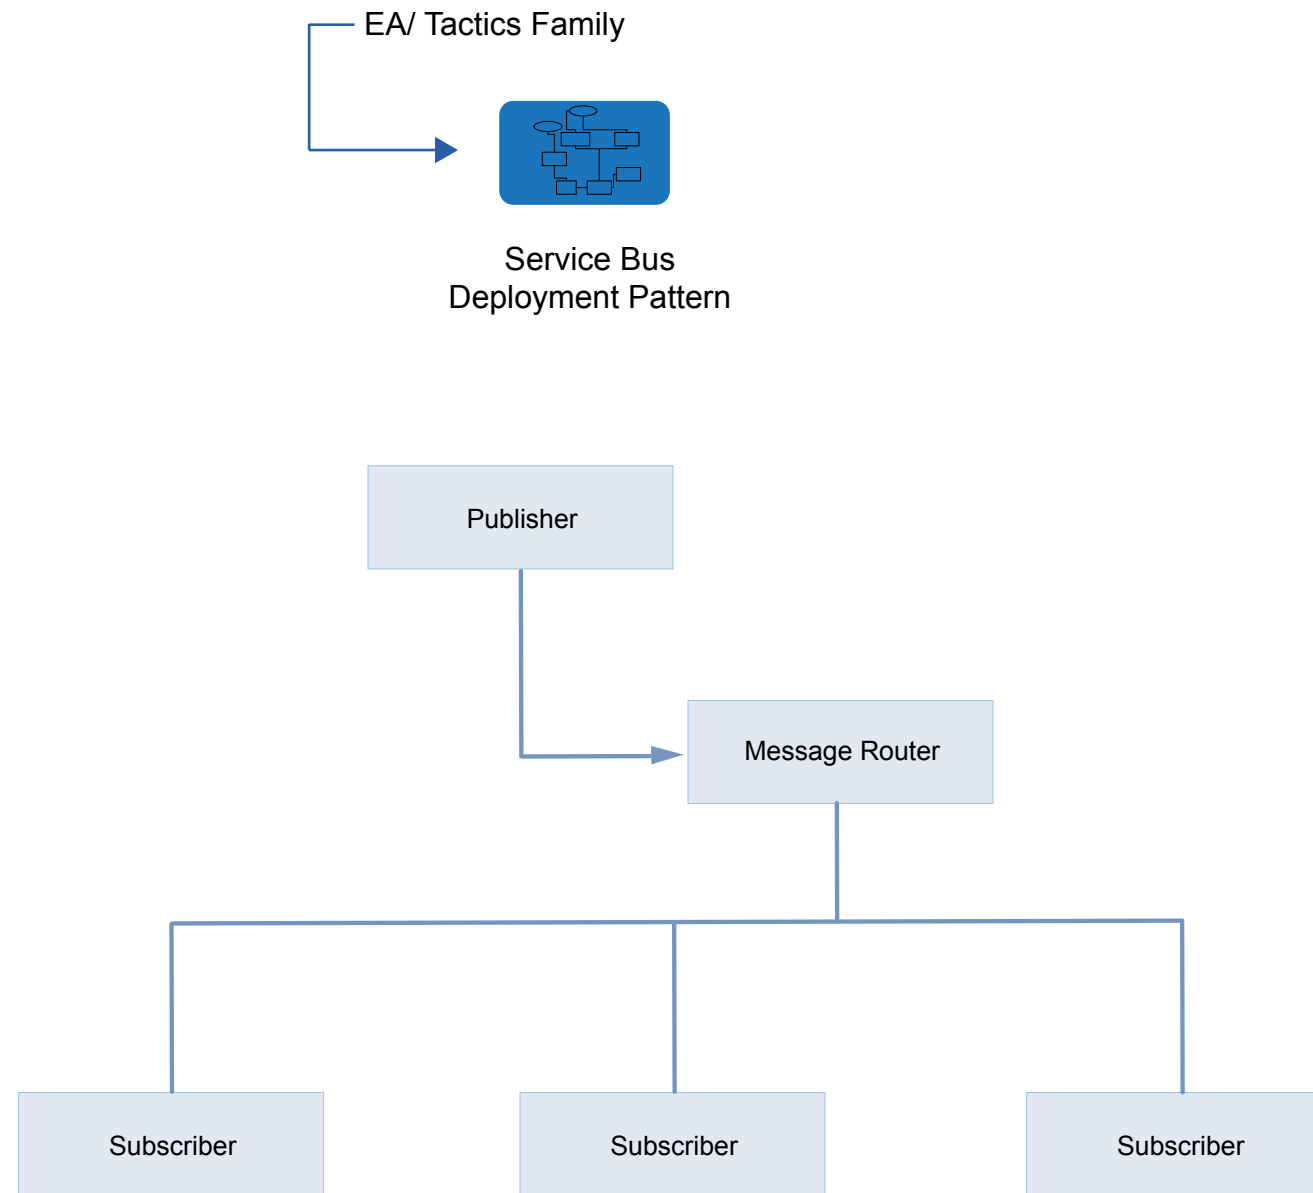


## Blueprint GPS

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



## 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

Functional Pattern Component	Traditional Deployment Pattern Name	Deployment Pattern Component	Guest Virtual Machine Consumption Characteristics			
			Compute (GHz)	Memory (GB)	Disk (GB)	Network (Gbps)
Business Logic	3 Tier Server	Application Server	4.25	8.60	384.62	6.51
Data Distribution Manager	Service Bus	Publishing Server	3.30	8.60	230.77	6.51
Data Source Communicator	3 Tier Server	Application Server	1.42	4.78	384.62	6.51
Edge Cache Repository	3 Tier Server	Web Server	2.36	6.69	384.62	6.51
Garbage Collector	3 Tier Server	Application Server	1.42	8.60	0.00	0.00
Media Repository	3 Tier Server	Database Server	4.25	8.60	692.31	6.51
Namespace Manager	Service Bus	Message Router Server	3.30	6.69	384.62	6.51
Non-Rich Media Transformation Engine	3 Tier Server	Application Server	4.25	8.60	384.62	6.51
Page Builder	3 Tier Server	Application Server	3.30	8.60	153.85	6.51
Personalization Repository	3 Tier Server	Database Server	2.36	8.60	538.46	6.51
Pub/Sub Manager	Service Bus	Message Router Server	3.30	8.60	538.46	6.51
Request Representation Mapper	3 Tier Server	Application Server	3.30	6.69	384.62	6.51
Rich Media Transformation Engine	3 Tier Server	Application Server	4.25	8.60	692.31	6.51
Service Request Manager	3 Tier Server	Web Server	3.30	8.60	230.77	6.51
Session Page Based Repository	3 Tier Server	Application Server	2.36	8.60	230.77	2.17
Targeted Source & Request Update Builder	3 Tier Server	Application Server	3.30	8.60	384.62	6.51

## 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.

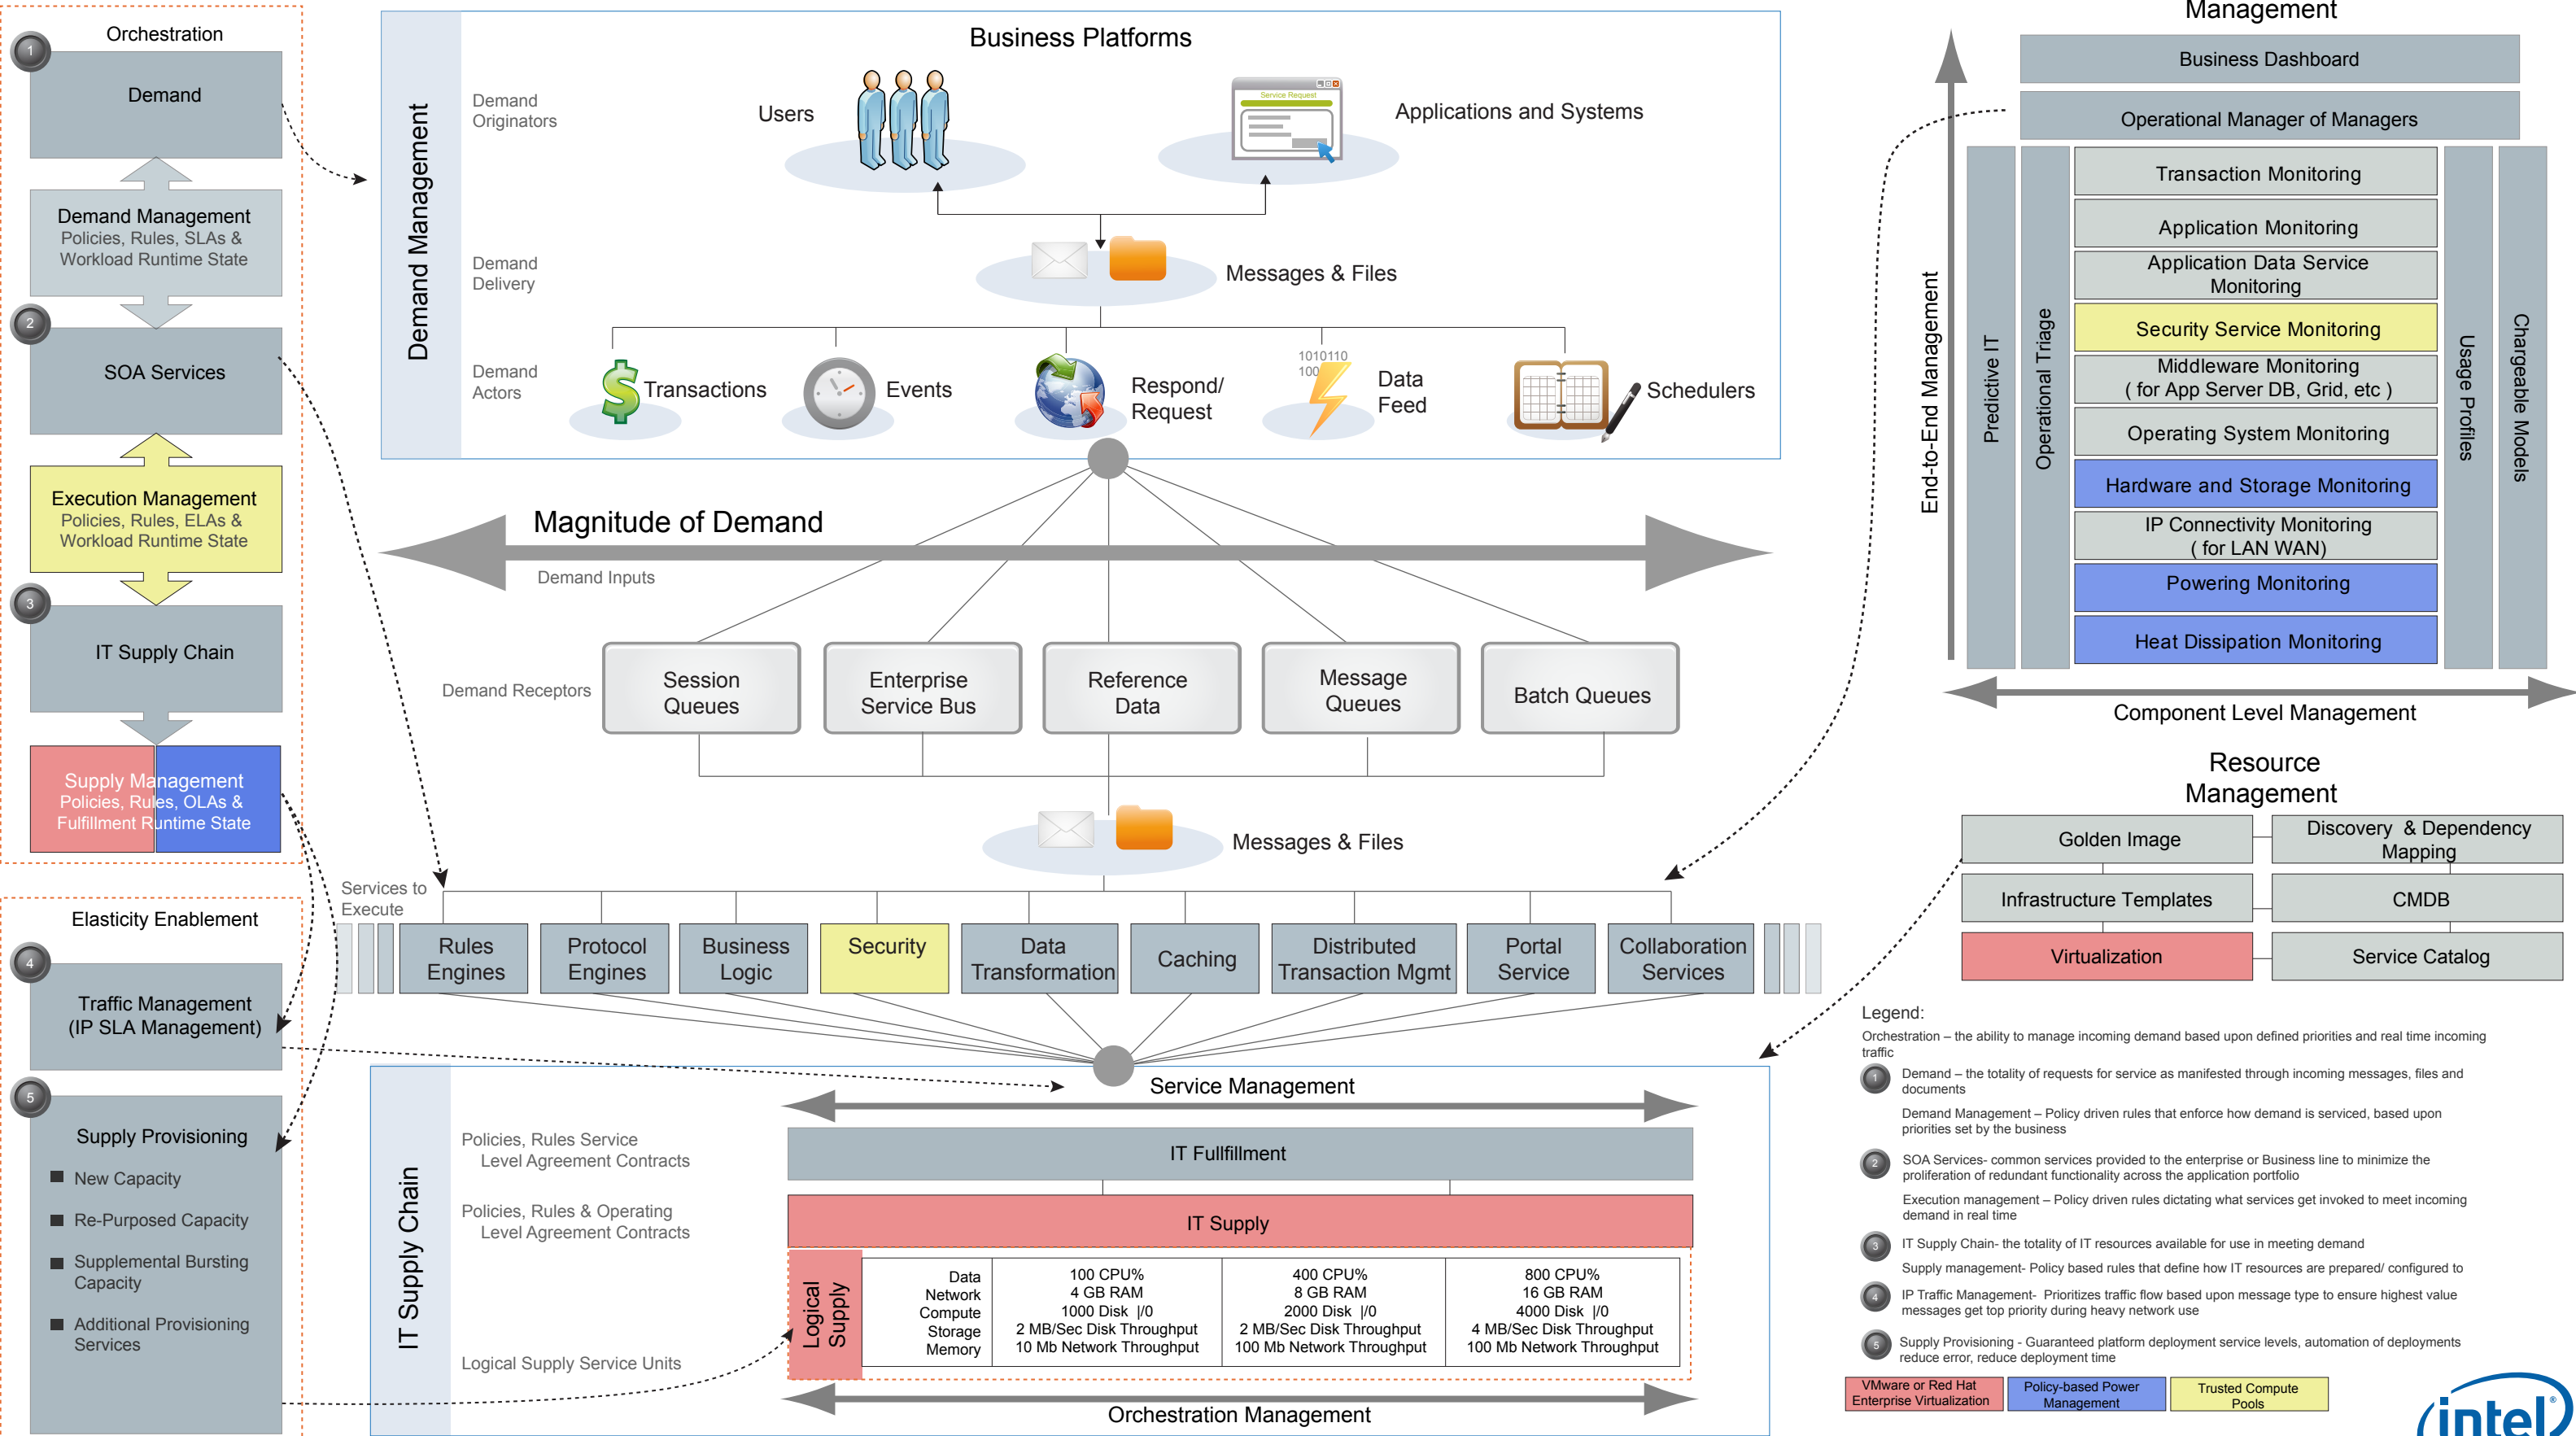


Capability Name	Vendor Name	Product Name	Model/Version	Quantity (#)	Base Attribute 1	Value	Unit	Base Attribute 2	Value	Unit	Base Attribute 3	Value	Unit
Compute Processor	Intel	Intel® Xeon®	processor X5570	1	Max Core Clock Speed	2.66	GHz	Max # of Cores	4	#	Max # of Threads	8	#
Compute Processor	Intel	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®	processor7555	0	Max Core Clock Speed	1.86	GHz	Max # of Cores	2	#	Max # of Threads	2	#
Compute Rack Server	Intel	Non-specific	Intel® S5500WB	6	Max # of Processor Sockets	2	#	Max # of Memory Slots	8	#	Max Memory Capacity	128	GB
Compute Memory	Non-specific	Non-specific	DDR3	44	Max Memory Capacity	4000	MB	Memory Type	DDR3	Categorical	Max # of Memory Ranks	2	#
Compute Chassis	Intel	Rack Chassis	Intel® SC5650BRP	1	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	6Tray-14 DDMs/Tray	4	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	104	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 # of 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 # of Systems Monitored	1024	#	Real-time Monitoring	Yes	Categorical	Availability Monitoring	Yes	Categorical
Golden Image Generation	VMware*	vSphere*	v4.1	1	Max # of Images Supported	Unlimited	#	Custom Image Templates	Yes	Categorical	Custom Media Repository	Yes	Categorical
Policy Based Provisioning	VMware*	vCenter*	v4.1	1	Agg. Physical Resources	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 Resource Balance	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 # of Virtual Switches	16	#	Max # of Virtual Ports	256	#	VLAN Tagging Support	Yes	Categorical
Virtual Machines	VMware*	vSphere*	v4.1	1	Max # of VMs per Server	512	#	Max Memory per VM	256	GB	Max CPU per VM	8	#
Virtual Machine Snapshots	VMware*	vSphere*	v4.1	1	Max # of Snapshots	Unlimited	#	Max # of Conc.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 App Support	1028	#	Virtual Media Support	Yes	Categorical	Virtual Catalog Support	Yes	Categorical
Cluster/Pool Balancing	VMware*	vSphere*	v4.1	1	Dynamic Resource Balance	Yes	Categorical	Resource Monitoring Sup.	Yes	Categorical	Power Mgmt Support	Yes	Categorical
Access Security Monitoring	VMware*	vShield*	v4.1	1	Policy Decision Point	Yes	Categorical	Policy Execution Point	Yes	Categorical	Policy Information Point	Yes	Categorical
Workload Orchestration & Management	VMware*	vCloud Director*	v1.0	1	Policy-based resource Mgmt	Yes	Categorical	Integrates Auto. Processes	Yes	Categorical	Exception Handling	Yes	Categorical
Security Monitoring	VMware*	vShield*	v4.1	1	Policy Decision Point	Yes	Categorical	Policy Execution Point	Yes	Categorical	Policy Information Point	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	64	bits
Compute Processor	Intel	Intel® Xeon®	processor X7560	6	Max Memory Size	144	GB	Bit Support	64	bits
Compute Processor	Intel	Intel® Xeon®	processor7555	0	Max Memory Size	144	GB	Bit Support	64	bits
Compute Rack Server	Intel	Non-specific	Intel® S5500WB	6	Max # of I/O Slots	6	#	Max I/O Bandwidth	40	Gbps
Compute Memory	Non-specific	Non-specific	DDR3	44	(null)	(null)	(null)	(null)	(null)	(null)
Compute Chassis	Intel	Rack Chassis	Intel® SC5650BRP	1	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	6Tray-14 DDMs/Tray	4	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	104	(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 HWReading	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 HWReading	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 HWReading	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
Access Security Monitoring	VMware*	vShield*	v4.1	1	Policy Access Point	Yes	Categorical	Single Sign-on	Yes	Categorical
Workload Orchestration & Management	VMware*	vCloud Director*	v1.0	1	Centralized Resource Mgmt	Yes	Categorical	Auditing	Yes	Categorical
Security Monitoring	VMware*	vShield*	v4.1	1	Policy Access Point	Yes	Categorical	Single Sign-on	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



- Legend:**
- Orchestration – the ability to manage incoming demand based upon defined priorities and real time incoming traffic
  - 1 Demand – the totality of requests for service as manifested through incoming messages, files and documents
  - Demand Management – Policy driven rules that enforce how demand is serviced, based upon priorities set by the business
  - 2 SOA Services- common services provided to the enterprise or Business line to minimize the proliferation of redundant functionality across the application portfolio
  - Execution management – Policy driven rules dictating what services get invoked to meet incoming demand in real time
  - 3 IT Supply Chain- the totality of IT resources available for use in meeting demand
  - Supply management- Policy based rules that define how IT resources are prepared/ configured to
  - 4 IP Traffic Management- Prioritizes traffic flow based upon message type to ensure highest value messages get top priority during heavy network use
  - 5 Supply Provisioning - Guaranteed platform deployment service levels, automation of deployments reduce error, reduce deployment time

VMware or Red Hat Enterprise Virtualization | Policy-based Power Management | Trusted Compute Pools



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.