# Intel® Cloud Builders: Nightly Batch File

This is a blueprint for a Nightly Batch File process doing many large transactions of Extract/Transform/Load (ETL) data processing, based on a Red Hat* Enterprise Virtualization Hypervisor 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 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 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 | 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, 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 |
| 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 | 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. |

# Intel® Cloud Builders: Nightly Batch File : Extract/Transform/Load Definition

## ETL - Cloud Deployed Components

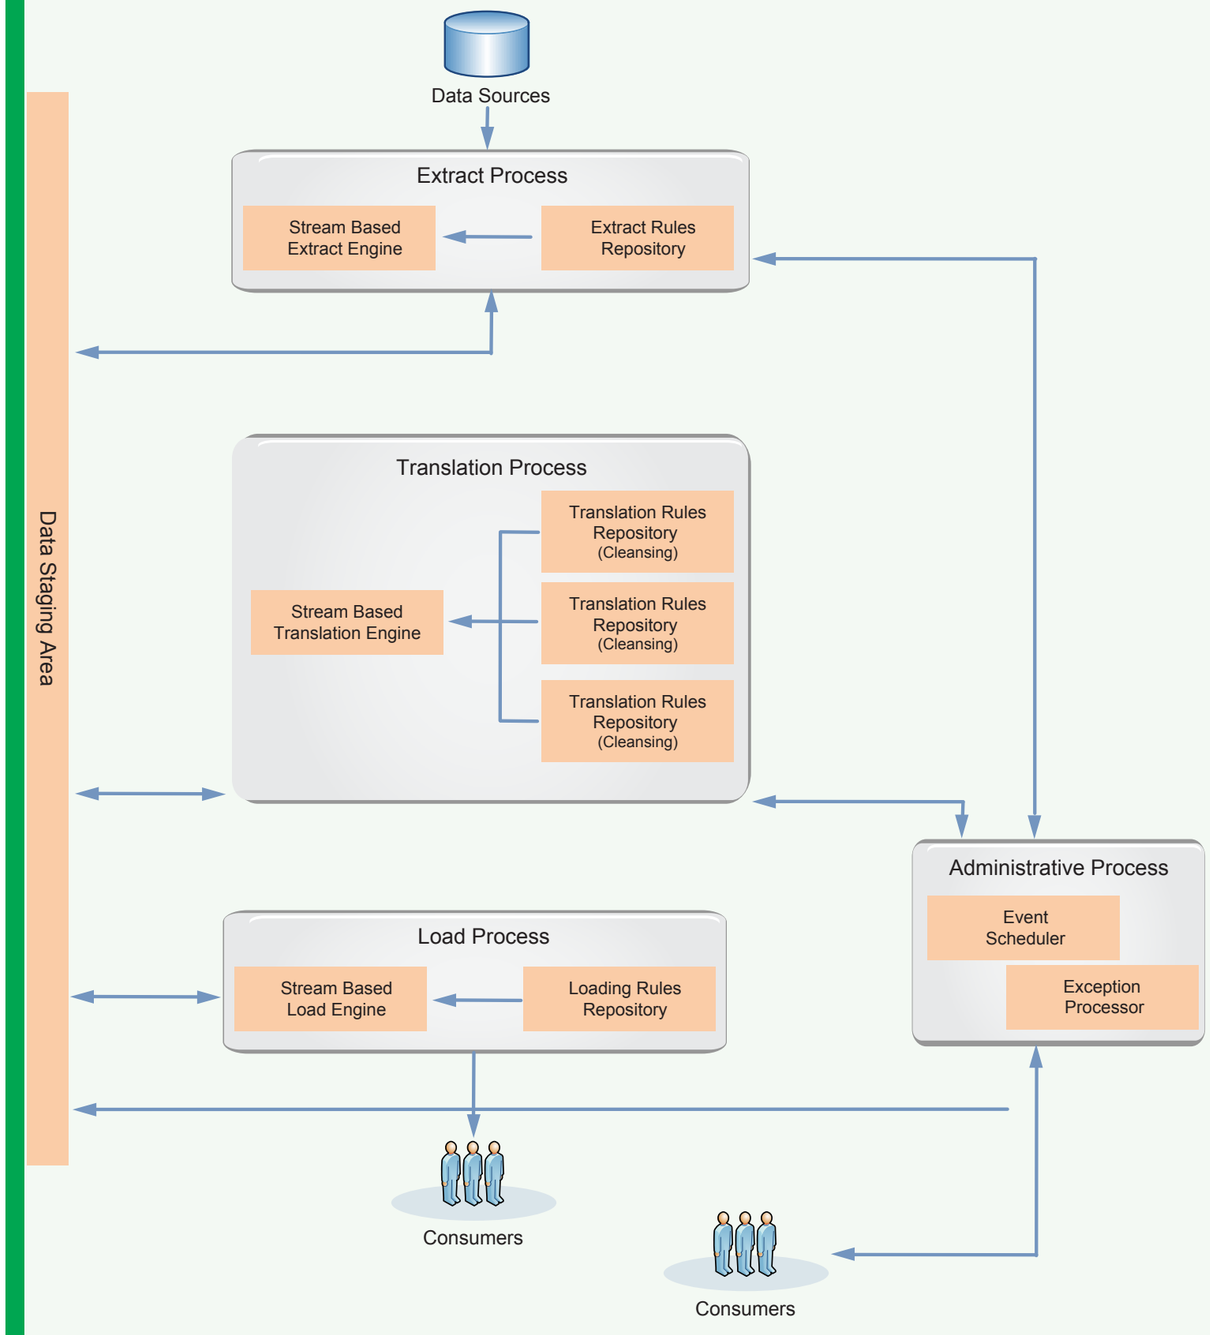


## BASELINE SELECTIONS

### Demand Characteristics

| | |
|---|---|
| Total # of Records (#) | Up to 100M |
| Runtime SLA (hours) | Less than 3 hrs |
| Total Data Load (TB) | Up to 250GB |
| Total Records Processed per Second (MB/s) | Up to 2MB/s |

### Resource Management Strategy

| | |
|---|---|
| Trusted Compute Pools: | Yes |
| Policy-based Power Management | Yes |

### Virtualization Management Strategy

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

## COMPUTE SELECTIONS

| | |
|---|---|
| Processor Chipset:<br>Intel® Xeon® processor X5570<br>Intel® Xeon® processor X7560<br>Intel® Xeon® processor L7555 | Quantity:<br>0<br>4<br>5 |
| Memory Modules:<br>DDR3 | Quantity:<br>13 |
| Network Adapter:<br>Intel® 82598EB 10 Gigabit Ethernet Controller | Quantity:<br>4 |
| Disk Type:<br>Disk Drive Module - SSD | Quantity:<br>19 |

## Pattern Function Descriptions

**Extract Rules Repository**- Contains rules that govern the extract process in a data driven manner.

**Event Scheduler**- Holds the operational schedule for all tasks and invokes processes as required.

**Data Staging Area**- Flexible storage pool that allows all 3 ETL processes to run simultaneously.

**Stream Based Translation Engine**- Transform multiple, arbitrarily sized and formatted data sets from one representation to another.

**Translation Rules Repository**- Rules for translating formats in a data driven, programmatic manner. Used because the rules change frequently.

**Stream Based Load Engine**- Distribute large data sets to multiple consumers, scheduled with variable delivery service levels.

**Exception Processor**- Provides a queue of all exceptions for the operator. Takes action to fix problems based upon operator commands.

**Stream Based Extract Engine**- Aggregate many sources of data into pre-configured information hierarchies, categories or record types.

## Blueprint GPS

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

# Intel® Cloud Builders: Nightly Batch File: Deployment Pattern for SOA: Pull

Illustrates a logical deployment architecture.

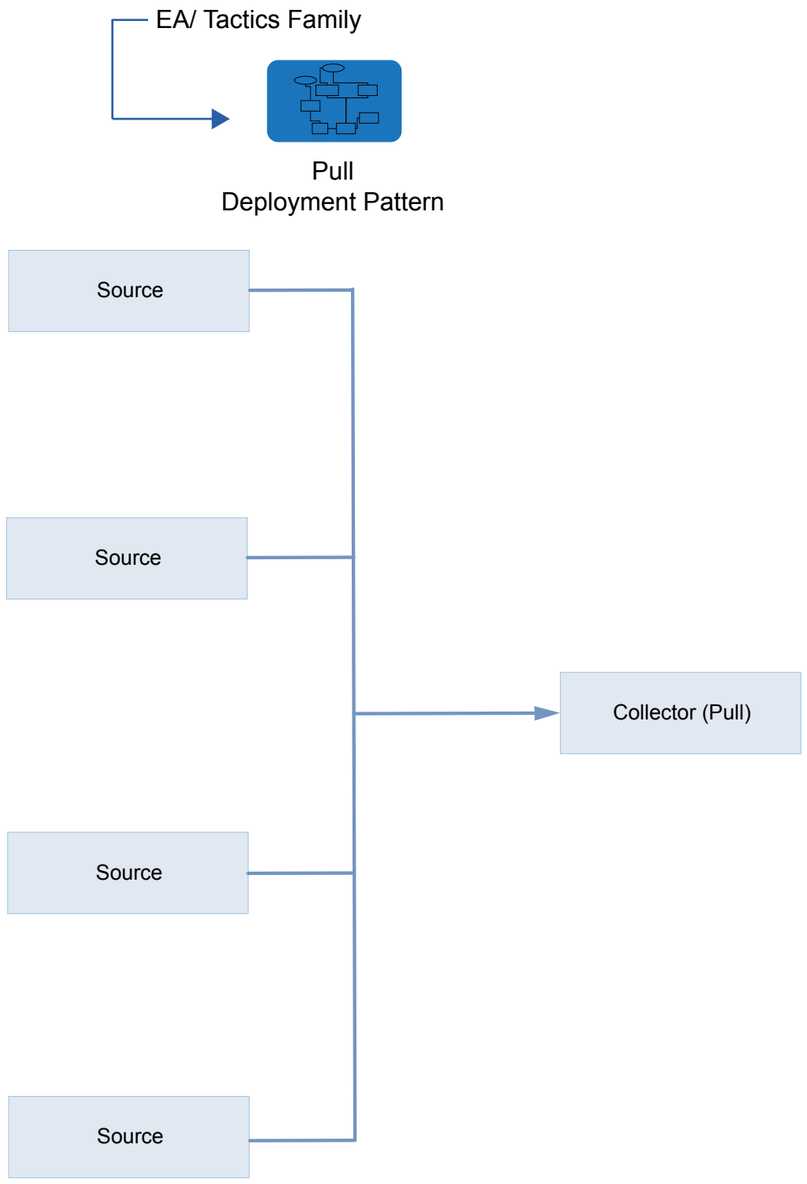


## Blueprint GPS

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

# Intel® Cloud Builders: Nightly Batch File: Deployment Pattern for SOA: Push

Illustrates a logical deployment architecture.

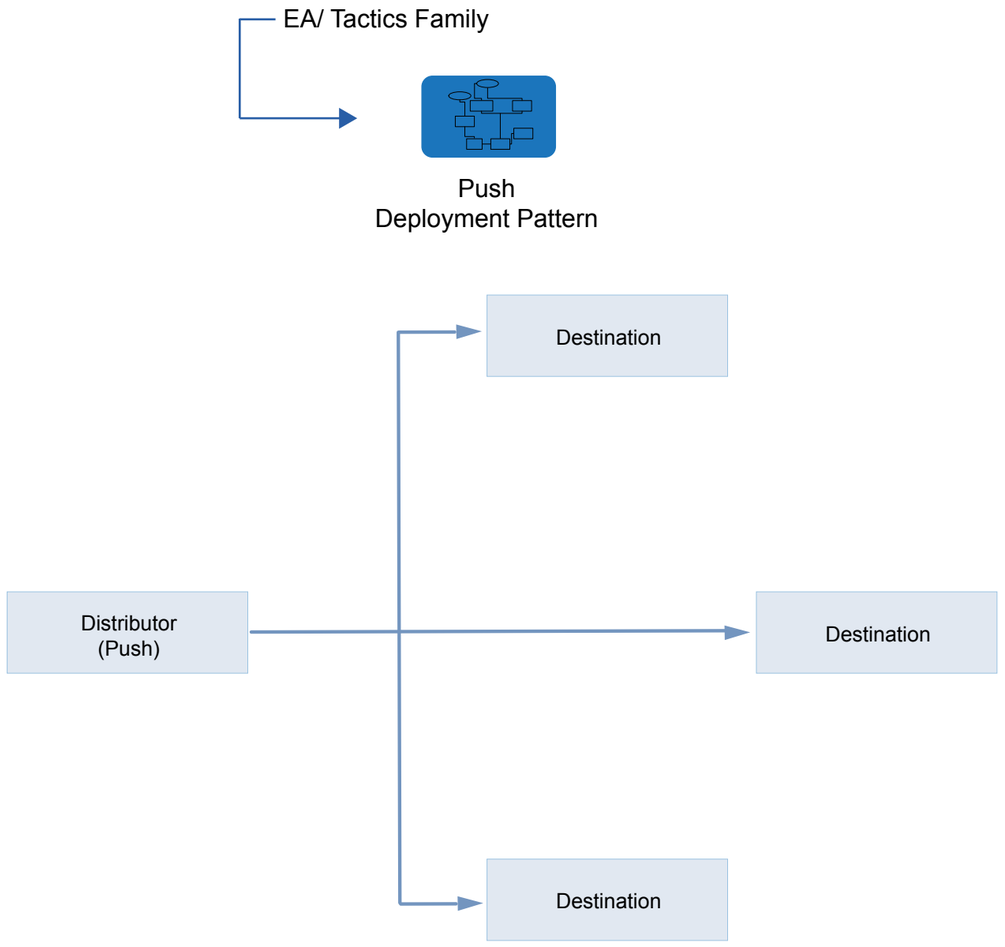


## Blueprint GPS

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

# Intel® Cloud Builders: Nightly Batch File: Deployment Pattern for SOA: 3 Tier Server

Illustrates a logical deployment architecture.

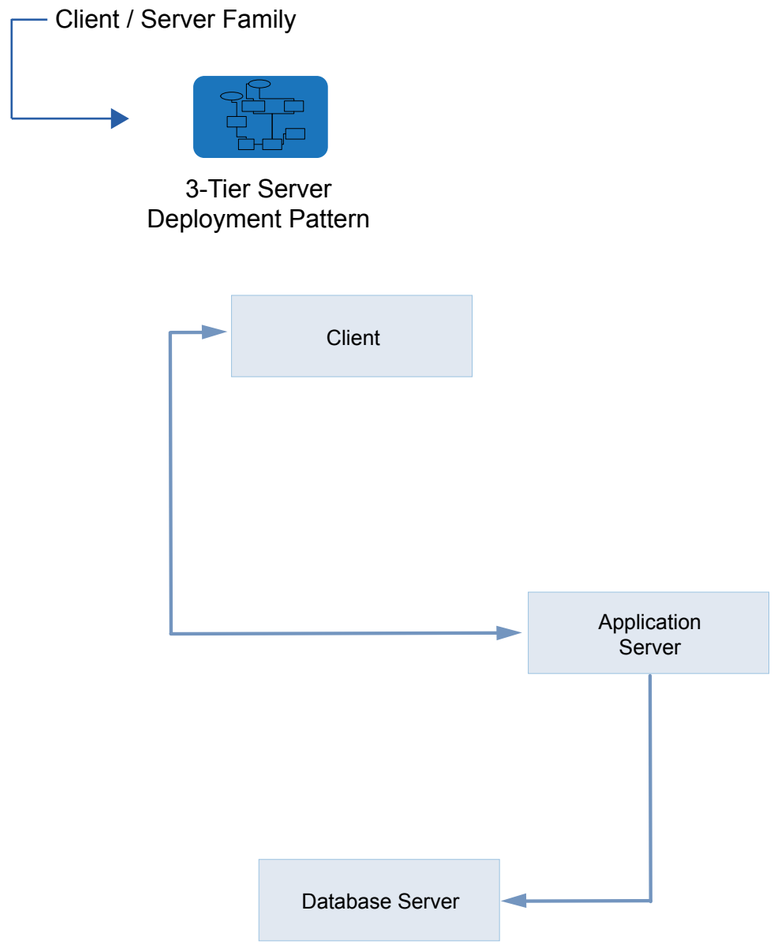


## Blueprint GPS

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

# Intel® Cloud Builders: Nightly Batch File: ETL Virtual Configuration Profile

| Functional Pattern Component | Traditonal Deployment Pattern Name | Deployment Pattern Component | Guest Virtual Machine Consumption Characteristics | | | |
|---|---|---|---|---|---|---|
| | | | Compute (GHz) | Memory (GB) | Disk (GB) | Network (Gbps) |
| Data Staging Area | 3 Tier Server | Database Server | 1.58 | 3.95 | 48.91 | 0.46 |
| Event Scheduler | 3 Tier Server | Web Server | 1.58 | 0.88 | 10.87 | 0.46 |
| Exception Processor | Pull | Web Server | 1.58 | 0.88 | 10.87 | 0.46 |
| Extract Rules Repository | Push | Collector Server | 1.58 | 1.32 | 10.87 | 0.46 |
| Loading Rule Repository | Pull | Distributor Server | 1.58 | 1.32 | 10.87 | 0.46 |
| Stream Based Extract Engine | Push | Collector Server | 3.96 | 3.95 | 38.04 | 0.46 |
| Stream Based Load Engine | 3 Tier Server | Distributor Server | 3.96 | 3.95 | 38.04 | 0.46 |
| Stream Based Translation Engine- Aligning | 3 Tier Server | Application Server | 7.13 | 3.95 | 16.30 | 0.46 |
| Stream Based Translation Engine- Cleansing | 3 Tier Server | Application Server | 7.13 | 3.95 | 16.30 | 0.46 |
| Stream Based Translation Engine- Summarizing | 3 Tier Server | Application Server | 7.13 | 3.95 | 16.30 | 0.46 |
| Translation Rules Repository- Aligning | 3 Tier Server | Application Server | 1.58 | 1.32 | 10.87 | 0.46 |
| Translation Rules Repository- Cleansing | 3 Tier Server | Application Server | 1.58 | 1.32 | 10.87 | 0.46 |
| Translation Rules Repository- Summarizing | 3 Tier Server | Application Server | 1.58 | 1.32 | 10.87 | 0.46 |

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: Nightly Batch File: 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 | 0 | Max Core Clock Speed | 2.66 | GHz | Max # of Cores | 4 | # | Max # of Threads | 8 | # |
| Compute Processor | Intel | Intel® Xeon® | processor X7560 | 4 | Max Core Clock Speed | 2.93 | GHz | Max # of Cores | 4 | # | Max # of Threads | 8 | # |
| Compute Processor | Intel | Intel® Xeon® | processor L7555 | 5 | Max Core Clock Speed | 1.86 | GHz | Max # of Cores | 2 | # | Max # of Threads | 2 | # |
| Compute Rack Server | Intel | Non-specific | Intel® S5500WB | 5 | Max # of Processor Sockets | 2 | # | Max # of Memory Slots | 8 | # | Max Memory Capacity | 128 | GB |
| Compute Memory | Non-specific | Non-specific | DDR3 | 13 | 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 | 4 | Max Throughput | 1 | Gbps | Max I/O | 8000 | IOPS | Max # of Ports | 4 | # |
| Converged Network Adapter | Intel | Eth Server Adapter | Intel® 82598EB | 0 | Max Throughput | 10 | Gbps | Max I/O | 40000 | IOPS | Max # of Ports | 2 | # |
| Storage Subsystem | Non-specific | Non-specific | 6Tray-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 | 19 | 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 | 0 | 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 | Red Hat* | RHEV | v2.2 | 1 | Max # of Images Supported | Unlimited | # | Custom Image Templates | Yes | Categorical | Custom Media Repository | Yes | Categorical |
| Policy Based Provisioning | Red Hat* | RHEV | v2.2 | 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 | Red Hat* | RHEV | v2.2 | 1 | Moniter Power Useage | 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 | Red Hat* | RHEV | v2.2 | 1 | Monitor Cluster State | Yes | Categorical | Synchronize Databases | No | Categorical | Create DBSnapshots | No | Categorical |
| VM Patch Management | Red Hat* | RHEV | v2.2 | 1 | OS Support | Windows | Categorical | VM Server Patch Mgmt | Yes | Categorical | VM Guest Patch Mgmt | Yes | Categorical |
| Dynamic Resource Pools | Red Hat* | RHEV | v2.2 | 1 | Dynamic Resource Balance | Yes | Categorical | Resource Monitoring Sup. | Yes | Categorical | Power Mgmt Support | Yes | Categorical |
| Thin Provisioning | Red Hat* | RHEV | v2.2 | 1 | Compute Thin Provisioning | Yes | Categorical | Memory Thin Provisioning | Yes | Categorical | Network Thin Provisioning | No | Categorical |
| Virtual Network Monitoring | Red Hat* | RHEV | v2.2 | 1 | Agent-less Monitoring | Yes | Categorical | Network Performance | Yes | Categorical | Dependency Mapping | No | Categorical |
| Virtual Storage Configuration Management | Red Hat* | RHEV | v2.2 | 1 | VM Guest Datastore Mgmt | Yes | Categorical | Storage Subsystem Control | No | Categorical | LUN Provisioning | Yes | Categorical |
| Virtual Disk Management | Red Hat* | RHEV | v2.2 | 1 | VM Guest Support | Windows | Categorical | VM Server Support | Local | Categorical | (null) | (null) | (null) |
| Virtual Machine Monitoring | Red Hat* | RHEV | v2.2 | 1 | Agent-less Monitoring | Yes | Categorical | VM State Monitoring | Yes | Categorical | (null) | (null) | (null) |
| Virtual Networks | Red Hat* | RHEV | v2.2 | 1 | Max # of Virtual Switches | 16 | # | Max # of Virtual Ports | 256 | # | VLAN Tagging Support | Yes | Categorical |
| Virtual Machines | Red Hat* | RHEV | v2.2 | 1 | Max # of VMs per Server | 512 | # | Max Memory per VM | 256 | GB | Max CPU per VM | 8 | # |
| Virtual Machine Snapshots | Red Hat* | RHEV | v2.2 | 1 | Max # of Snapshots | Unlimited | # | Max # of Conc.Guest Builds | 1 | # | (null) | (null) | (null) |
| Virtual Resource Monitoring | Red Hat* | RHEV | v2.2 | 1 | Agent-less Monitoring | Yes | Categorical | VM Compute Monitoring | Yes | Categorical | VM Storage Monitoring | Yes | Categorical |
| Cloud Self-Service Portal | Red Hat* | RHEV | v2.2 | 1 | Max Virtual App Support | 1028 | # | Virtual Media Support | Yes | Categorical | Virtual Catalog Support | Yes | Categorical |
| Cluster/Pool Balancing | Red Hat* | RHEV | v2.2 | 1 | Dynamic Resource Balance | Yes | Categorical | Resource Monitoring Sup. | Yes | Categorical | Power Mgmt Support | Yes | Categorical |
| Access Security Monitoring | Red Hat* | RHEL | v5.5 | 1 | Policy Decision Point | Yes | Categorical | Policy Execution Point | Yes | Categorical | Policy Information Point | Yes | Categorical |
| Workload Orchestration & Management | Red Hat* | RHEV | v2.2 | 1 | Policy-based resource Mgmt | Yes | Categorical | Integrates Auto. Processes | Yes | Categorical | Exception Handling | Yes | Categorical |
| Security Monitoring | Red Hat* | RHEL | v5.5 | 1 | Policy Decision Point | Yes | Categorical | Policy Execution Point | Yes | Categorical | Policy Information Point | Yes | Categorical |

# Intel® Cloud Builders: Nightly Batch File: Logical Configuration Page 2

| 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 | 0 | Max Memory Size | 144 | GB | Bit Support | 64 | bits |
| Compute Processor | Intel | Intel® Xeon® | processor X7560 | 4 | Max Memory Size | 144 | GB | Bit Support | 64 | bits |
| Compute Processor | Intel | Intel® Xeon® | processor L7555 | 5 | Max Memory Size | 144 | GB | Bit Support | 64 | bits |
| Compute Rack Server | Intel | Non-specific | Intel® S5500WB | 5 | Max # of I/O Slots | 6 | # | Max I/O Bandwidth | 40 | Gbps |
| Compute Memory | Non-specific | Non-specific | DDR3 | 13 | (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 | 4 | (null) | (null) | (null) | (null) | (null) | (null) |
| Converged Network Adapter | Intel | Eth Server Adapter | Intel® 82598EB | 0 | (null) | (null) | (null) | (null) | (null) | (null) |
| Storage Subsystem | Non-specific | Non-specific | 6Tray-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 | 19 | (null) | (null) | (null) | (null) | (null) | (null) |
| Disk Drive Module | Non-specific | Non-specific | DD Module - SSD | 0 | (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 | Red Hat* | RHEV | v2.2 | 1 | Media Transcription | Yes | Categorical | Automatic Patch Insertion | Yes | Categorical |
| Policy Based Provisioning | Red Hat* | RHEV | v2.2 | 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 | Red Hat* | RHEV | v2.2 | 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 | Red Hat* | RHEV | v2.2 | 1 | Log Shipping | Yes | Categorical | Create Database Mirror | Yes | Categorical |
| VM Patch Management | Red Hat* | RHEV | v2.2 | 1 | (null) | (null) | (null) | (null) | (null) | (null) |
| Dynamic Resource Pools | Red Hat* | RHEV | v2.2 | 1 | (null) | (null) | (null) | (null) | (null) | (null) |
| Thin Provisioning | Red Hat* | RHEV | v2.2 | 1 | Datastore Thin Provisioning | Yes | Categorical | (null) | (null) | (null) |
| Virtual Network Monitoring | Red Hat* | RHEV | v2.2 | 1 | (null) | (null) | (null) | (null) | (null) | (null) |
| Virtual Storage Configuration Management | Red Hat* | RHEV | v2.2 | 1 | (null) | (null) | (null) | (null) | (null) | (null) |
| Virtual Disk Management | Red Hat* | RHEV | v2.2 | 1 | (null) | (null) | (null) | (null) | (null) | (null) |
| Virtual Machine Monitoring | Red Hat* | RHEV | v2.2 | 1 | (null) | (null) | (null) | (null) | (null) | (null) |
| Virtual Networks | Red Hat* | RHEV | v2.2 | 1 | (null) | (null) | (null) | (null) | (null) | (null) |
| Virtual Machines | Red Hat* | RHEV | v2.2 | 1 | Max I/O per VM | (null) | IOPS | Max Throughput per VM | (null) | Gbps |
| Virtual Machine Snapshots | Red Hat* | RHEV | v2.2 | 1 | (null) | (null) | (null) | (null) | (null) | (null) |
| Virtual Resource Monitoring | Red Hat* | RHEV | v2.2 | 1 | VM Memory Monitoring | Yes | Categorical | (null) | (null) | (null) |
| Cloud Self-Service Portal | Red Hat* | RHEV | v2.2 | 1 | Virtual Organization Support | Yes | Categorical | Partitioned Network Sup. | Yes | Categorical |
| Cluster/Pool Balancing | Red Hat* | RHEV | v2.2 | 1 | Automatic Pool Balancing | Yes | Categorical | Policy Based Pool Balance | Yes | Categorical |
| Access Security Monitoring | Red Hat* | RHEL | v5.5 | 1 | Policy Access Point | Yes | Categorical | Single Sign-on | Yes | Categorical |
| Workload Orchestration & Management | Red Hat* | RHEV | v2.2 | 1 | Centralized Resource Mgmt | Yes | Categorical | Auditing | Yes | Categorical |
| Security Monitoring | Red Hat* | RHEL | v5.5 | 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

Orchestration

1 Demand

Demand Management
Policies, Rules, SLAs & Workload Runtime State

2 SOA Services

Execution Management
Policies, Rules, ELAs & Workload Runtime State

3 IT Supply Chain

Supply Management
Policies, Rules, OLAs & Fulfillment Runtime State

Elasticity Enablement

4 Traffic Management (IP SLA Management)

5 Supply Provisioning
- New Capacity
- Re-Purposed Capacity
- Supplemental Bursting Capacity
- Additional Provisioning Services

Demand Management

Business Platforms

Demand Originators: Users, Applications and Systems

Demand Delivery: Messages & Files

Demand Actors: Transactions, Events, Respond/ Request, Data Feed, Schedulers

Magnitude of Demand

Demand Inputs

Demand Receptors: Session Queues, Enterprise Service Bus, Reference Data, Message Queues, Batch Queues

Messages & Files

Services to Execute: Rules Engines, Protocol Engines, Business Logic, Security, Data Transformation, Caching, Distributed Transaction Mgmt, Portal Service, Collaboration Services

IT Supply Chain

Service Management

Policies, Rules Service Level Agreement Contracts

IT Fullfillment

Policies, Rules & Operating Level Agreement Contracts

IT Supply

Logical Supply Service Units

Logical Supply

| | | | |
|---|---|---|---|
| Data<br>Network<br>Compute<br>Storage<br>Memory | 100 CPU%<br>4 GB RAM<br>1000 Disk \|/0<br>2 MB/Sec Disk Throughput<br>10 Mb Network Throughput | 400 CPU%<br>8 GB RAM<br>2000 Disk \|/0<br>2 MB/Sec Disk Throughput<br>100 Mb Network Throughput | 800 CPU%<br>16 GB RAM<br>4000 Disk \|/0<br>4 MB/Sec Disk Throughput<br>100 Mb Network Throughput |

Orchestration Management

Operations Infrastructure Lifecycle Management

End-to-End Management

Business Dashboard

Operational Manager of Managers

Predictive IT

Operational Triage
- Transaction Monitoring
- Application Monitoring
- Application Data Service Monitoring
- Security Service Monitoring
- Middleware Monitoring ( for App Server DB, Grid, etc )
- Operating System Monitoring
- Hardware and Storage Monitoring
- IP Connectivity Monitoring ( for LAN WAN)
- Powering Monitoring
- Heat Dissipation Monitoring

Usage Profiles

Chargeable Models

Component Level Management

Resource Management

| | |
|---|---|
| Golden Image | Discovery & Dependency Mapping |
| Infrastructure Templates | CMDB |
| Virtualization | Service Catalog |

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