HTNG Above Property Systems Considerations Version 1.0

Publication Date
7 January 2016

About HTNG

Hotel Technology Next Generation (HTNG) is a non-profit association with a mission to foster, through collaboration and partnership, the development of next-generation systems and solutions that will enable hoteliers and their technology vendors to do business globally in the 21st century. HTNG is recognized as the leading voice of the global hotel community, articulating the technology requirements of hotel companies of all sizes to the vendor community. HTNG facilitate the development of technology models for hospitality that will foster innovation, improve the guest experience, increase the effectiveness and efficiency of hotels, and create a healthy ecosystem of technology suppliers.

Table of contents

Document Information

1.1 Document History

Version	Date	Author	Comments
0.01	2014-12-01	Above Property Systems (APS) Workgroup	Initial Outline Drafted
0.20	2015-02-21	Thomas Nievergelt	Business Considerations, Identity Management Drafted
0.30	2015-04-01	Brian Alessi	Financial Considerations Drafted
0.35	2015-05-11	Steve D'Erasmo	Other Considerations Updated
0.40	2015-06-15	John Bell	Updates to Financial Considerations, Other Considerations
0.50	2015-07-03	Eric Sullender	Caching, VPNs, CDNs, Monitoring and Availability, Training, Public and Private Webservices Drafted
0.60	2015-09-15	Mark Haley	Updates to Business Considerations
0.75	2015-10-01	Vijay Raghavan	Critical Service Levels Drafted
0.90	2015-11-05	APS Workgroup	Numerious Updates, Comments and Edits
0.95	2015-12-01	Lo Li, Brian Alessi, Thomas Nievergelt, Patrick Dunphy	Numerous Updates to Other Considerations
0.98	2015-12-10	APS Workgroup	Formal Comments into Final Draft
0.99	2015-12-20	HTNG Staff	Copy Edits and Merged to HTNG Style Guidelines
1.0	2016-01-07	APS Workgroup	Formal Publication to HTNG Members

1.2 Overview

On-Premises vs Cloud Computing: Which one is better? Which model makes more sense for your organization and your business needs? Which will serve the organization best as you grow? There is not a clearly right or wrong answer. One size will not fit all as you embark upon the task of weighing your options. Business value and the ability of your organization to leverage a true enterprise computing environment may well outweigh any additional costs you might incur. Total Cost of Ownership (TCO) will be an important deciding factor in any decision. Ensure you are considering all costs, whether they are from the application vendor partner, from the data center provider, from an Internet service provider, from your own organizational structure or from the property itself. What factors are in place today that can be leveraged to offset any "cloud computing" costs? You may be already paying for many of the pre-requisite services, making your TCO lower than you may have thought.

Finally, look at the intangibles. These factors are equally important in any decision making process. How will the operation be affected? Will your organization see added benefits by having all of your data in an enterprise environment? How will a "cloud" solution work to support the various geographies where your organization resides? What about the support and disaster recovery models? Will it be sustainable? Is "time to market" a compelling argument? While there is a myriad of factors to consider when making such a decision, knowing many of the key major factors to consider when evaluating the merits of an on-premises and a cloud-based software solution will make the task more manageable and will allow you to move through these topics in a methodical fashion.

2 Business Considerations

2.1 Introduction

This document summarizes above property solutions business considerations. The workgroup focused less on technical topics, and more on important criteria for the hotelier. You will find references to detailed documentation within the following content.

2.2 Business Value

When comparing an on-premises solution against a hosted, or cloud[1] solution, it is very common that a property is comparing the same exact software solution from the same vendor partner. The only tangible difference is where the actual application data will reside; on a server device at the property or on a server device in a hosted data center facility away from the property.

In most cases, the functionality of the software will be exactly the same in both models. The vendor will offer the same software version and set of features in either deployment model. It is important for the property to verify this situation to ensure that added features are not available in a hosted scenario or, more importantly, features the hotel may already enjoy and rely upon today in their premises-based installation are not eliminated or restricted when that same application moves to a hosted platform. For example, certain interfaces or data extracts that are readily enabled in an on-premises implementation can be compromised or forbidden by the vendor in a hosted scenario.

From a pure business value perspective, the feature/functionality set provided by the application provider in each deployment model should not add or detract from the overall business value. However, there are some key points to highlight to ensure the property comes into the evaluation of an on-premises versus a cloud-based software solution with a comprehensive understanding.

2.3 Ability to Leverage an Enterprise Database

In a hosted environment, economies of scale can be achieved by not requiring each individual property to purchase and maintain dedicated server hardware for use by both the application software and by the application database. Many application providers who offer a “cloud” solution will provide flexibility for all of your properties to share server hardware while, at the same time, separating the data from one property to the next and ensuring data privacy. This will become a critical consideration for organizations who work in a franchisee environment.

[1] Defining the distinction between “hosted” and “cloud” solutions is meaningful and relevant, but beyond the scope of this document. Recognize that the terms are not interchangeable and are often misused.

Through individual property setup of database management rules, it is quite possible for most applications to be configured where the organizational guidelines can be established and enforced while, at the same time, protecting one property's sensitive data (sales, profiles, customers, sales leads, etc.) from other franchisee data. This can be done while still housing all properties' data in the same overall database structure. When considering a hosted application environment, it is important for the organization to determine to what level of protection each individual property's data must be protected.

Once enterprise database structure rules have been established, a foundation will exist for organizations to leverage the fact that all of their property data now resides in one physical location (the hosting center) as opposed to residing on individual servers at each property. From the enterprise, or data center level, some interesting opportunities will then present themselves.

2.4 Time to Market

Business needs and objectives evolve over time. New requests are made on a regular basis to change or enhance existing functionality across similar systems deployed throughout your company's estate. While individual upgrades are always a possibility to consider, each one will require planning, operational downtime and skilled resources to complete. Larger companies with hundreds or thousands of locations are sometimes only able to complete one meaningful software upgrade per year as it takes dedicated resources that long to deploy the change across the entire estate.

In this model, it is common to have the business come and make new functionality requests of your IT organization long before you have had the opportunity to fully deploy the original update across all properties in your estate. From a business perspective, they are likely not too keen to hear of any limitations to deploy updates the minute they are ready. The exercise becomes a bit of a treadmill as the IT team may likely feel they are never able to complete an upgrade cycle prior to the next request piling up from their business partners.

A hosted deployment model does provide some relief in this area. Through the use of enterprise databases, your entire estate of properties – or, at a minimum, clusters of your properties – can receive the same software updates concurrently. This has the potential to dramatically decrease your "time to market" when deploying application updates and changes. While each update still requires the same level of planning and operational downtime, one set of activities creates the opportunity to touch a multitude of properties at the same time and shorten your overall total "time to market."

2.5 Effects of Selling a Property

Let's assume that you have made the decision to move a number of your locations to an enterprise, hosted model. The next logical scenario to define is "what happens when one of your properties leaves the organization?"

Does the enterprise model provide the flexibility for the data associated with that one property to be surgically extracted from the enterprise database and given to the property that is leaving the group? Is all of that property's data rightfully theirs or does a subset belong to your organization? In a property management system scenario for example, does the property that is leaving have rights to all reservation data and historical sales data? What about group business and future leads? Does this information belong to your organization or to the entity that is leaving you? Does the application provider have a mechanism to extract one property's data from the enterprise database or merely make a copy of that property's data while still leaving the offending property's data behind to clutter up your enterprise database?

The inverse scenario will also present itself eventually. Once the enterprise database has been established, how easy will it be for the vendor to add a new hotel property's data structure to your enterprise should your organization opt to add a new hotel to the portfolio? Knowing your partner's capabilities (up front) will make ensure a smooth transition for this type of scenario (and many more). Having a firm understanding of what is possible technically and what must be accounted for legally will save your organization quite a bit of time on the backend.

2.6 Information Ownership

If each hotel in a group of hotels is responsible for the data they need, then the responsibility for and ownership of the data often clearly falls with the individual property. In this case, if the property leaves the group, the information leaves with the property.

However, if the data is housed above property, the responsibility and ownership may fall on different parties. For example, a central reservation system for a chain of franchised hotels may contain availability, customer and reservation data in the same data store. Individual hotels may have access to the data that is specific for the hotel, but unless provisions and agreements have been made, the data may not leave if the hotel decides to leave the group. In a Software as a Service (SaaS) model, the data is contained within the application but the application is owned and operated by a third party. In this situation, the provider may potentially claim ownership of the data particularly when in the midst of a larger dispute. Hoteliers must ensure that hosting and SaaS contracts clearly delineate data ownership. Contract terms must also define vendor obligations to support exporting, copying or otherwise relocating data owned by the hotel or customer.

2.7 Risk

Risk management changes in a cloud or hosted environment compared to risk in many on-premises installations. Risk of a data breach, for example, is perhaps greater in a multi-property hosted scenario, where there is data about many more customers to be stolen than in any single site. This risk factor may be offset by the presumed greater expertise in system administration, security and monitoring expected in a professional data center operation.

Credit card data security becomes a challenge in a hosted environment, where PCI requirements explicitly establish that the merchant (i.e. the hotel) is liable for a breach. But in a hosted scenario, the merchant has little control or impact over the administration of the system, including security, monitoring and testing. The only impact the merchant can have is before signing an agreement, with contractual requirements for industry-standard security measures on the part of the vendor and the maximum possible indemnification negotiated.

2.8 Further Considerations

Consideration	Summary
Vendor Management	Utilizing cloud services may take more resources to manage the vendor (or service), rather than the more technical issues. Often, ensuring vendor adherence to SLAs and problem resolution need to be considered. Defining the key metrics and verification of those metrics is important. “Trust but verify.”
Data Ownership	Changing business affiliation (franchises) has different consequences when utilizing cloud services. If you are contemplating above property, you must maintain and contractually document ownership of your data in case it needs to be moved elsewhere.
Change Management	• Version control needs to be built into the system. • Staff and vendor roles must be clearly defined. • Regardless of who is responsible for the execution of the change, when you go above property, how changes occur is different. Pace of change, intervention, seamless changes, and risk are all factors that are affected by change management in the cloud. Keeping track of who is responsible for what is important. Keeping stakeholders involved is important for acceptance of change. Because the speed of change is quickened, keeping users trained and informed becomes more important (as well as who is responsible for supporting and communicating the change).
Data Recovery	Third party tools exist to extract data from cloud installations, but this will be complicated by the existing tools and methods available by your cloud service provider. Further, extracting data from a cloud provider in a franchise environment (such as when a physical hotel changes flags or brands) needs to be thoroughly reviewed and understood as part of the selection process. Disaster recovery scenarios (including data recovery) may be seamless to the end user in SaaS deployments, but may

Consideration	Summary
	be under the customer's responsibilities to plan and execute in PaaS and IaaS deployments.
Identity Management	Identity Management is an enterprise's strategy to control and track its users' systems access. The approach increases security and efficiency considerably. Different user account directory integration technologies can be applied.
Service Level Agreements (SLA)	Metrics of uptime, performance and data integrity as part of an SLA must be clearly defined. Most cloud vendors will not support the last mile internet connectivity. That is generally supported by the hotel or customer. You need to define what your business critical operations are; and if 100% availability is not available, offline modes are needed.
Support Criteria	Support criteria should be laid out in any SLA, legal or contracting document.
Monitoring	New points of monitoring are available in the cloud. SLAs often dictate what monitoring points you have.
Compliance	Verifying compliance across 100 physical installations vs a cloud offering is a much different undertaking, and can be very simplified. Cloud services allow access to metrics that were not previously readily available. Additionally, new tools are available to audit compliance when utilizing cloud architectures and vendors. Cloud system consumers should be aware that compliance issues, culpability and risk are complicated and some of that shifts to the cloud vendor.
Legal / Contracting	Paying close attention to legal and contractual terms is vitally important.
Financial	Cloud solutions often shift budget resources from a Cap-Ex model to an Op-Ex model, where upfront costs are dramatically reduced and longer term costs rise. A shift in resource focus (e.g., from infrastructure engineers to cloud developers) is also an important consideration to take into account. Ultimately, a thorough analysis of TCO over a long term period is required to gauge the financial changes that may result from a cloud deployment. Detailed information is available later in this document.
Transition	Transitions include downtime when switching to a cloud system, as well as potential data loss (for which you and the vendor need a disaster recovery plan).
Performance	Network latency of applications may be an issue. High performance applications can be configured to expand on an as

Consideration	Summary
	needed or automatic basis for necessary resources (CPU, RAM, disk space etc.). This may have an effect on costs, based on your particular vendor or situation.
Scalability	X as a service generally dictates your scalability constraints. (Platform, infrastructure, software, etc.) Dedicated hosting environments, for example, are not necessarily inherently scalable. SaaS offerings are typically inherently scalable, but there may be cost implications.
APIs	Application interfaces are indispensable when operating multi-solution environments. Providers need to deliver standard interfaces adhering to latest methodology. They need to have proven records of cooperation with third-party providers operating within the hospitality industry.
Reporting	Third-party tools exist to extract data from cloud installations, but this will be complicated by the existing tools and methods available by your cloud service provider. Further, extracting data from a cloud provider in a franchise environment (such as when a physical hotel changes flags or brands) needs to be thoroughly reviewed and understood as part of the selection process. Disaster recovery scenarios (including data recovery) may be seamless to the end user in SaaS deployments, but may be under the customer's responsibilities to plan and execute in PaaS and IaaS deployments.
Training	You need to be aware of who is responsible for training, and build it into your contracts. Continuous training and follow up are often overlooked in terms time and cost.
Availability	Availability is generally expressed in up-time (variances of 99% are typical). Measurements are typically taken on a monthly basis, but may vary by vendor or cloud service. There are potential advantages to shifting to a cloud service, where a vendor may have much more experience and capability to provide availability than an on-premises deployment, but a constant and stable internet connection may be an issue in certain parts of the world.
Architecture & Design	Architecture designs of a SaaS, PaaS and IaaS cloud deployment vary widely. Scalability, availability, performance, data recovery, and identity management are just some of the considerations you should examine when creating your own applications.
Security	Security models must include PCI compliance, appropriate modeling (especially in the case of tenant environments),

Consideration	Summary
	encryption of PII data, and state of the art intrusion prevention measures. Cloud deployments are often an extension of your environment, where the same processes and procedures should be deployed and enforced as on property. This often involves a contractual agreement, but must be verified. As opposed to on-property deployments, cloud deployments often include more than one point demarcation, so multiple security enforcement activities should be considered. Vendors should have a regular cadence of security scans to identify vulnerabilities in their applications.
Hosting	Private and public clouds each have advantages, depending on business needs and requirements. Despite these differences, data center requirements (and certifications) apply in both circumstances. Response times are a major consideration when selecting or evaluating a hosting provider.

3 Financial Considerations

Above property services generally change your financial model. Typical on-premises systems have significant capital investment. These costs include (but are not limited to):

- Hardware
- Software / Licenses
- Hosting
- Support
- Transaction
- Upgrades
- Maintenance
- Implementation Services

Generally speaking, these costs are transformed from a capital-heavy total cost of ownership to an expense-heavy total cost of ownership. If your organization operates a P&L, it can be challenging to justify a change in the financial model. As a general rule, a five-year TCO is accepted. However, due to hardware depreciation (which may or may not need capital re-investment after five years), sometimes a 6- or even 7-year model is used. It depends on who suffers the expense for capital.

As an example, under a SaaS model, an organization may reduce its hardware costs and its software, license, hosting, support, upgrade and maintenance costs will be transformed into a commodity-based service cost. This is commonly referred to as a subscription fee. These fees are operating expenses rather than capital.

In all the below scenarios we will be referencing a Food and Beverage Point of Sale system for a hotel/location with 350 physical Point of Sale (POS) terminals.

In all scenarios, each individual circumstance will be different and is not intended to be representative of actual costs or savings. However, it's an exercise your hotel/company should go through when evaluating on- property vs. some sort of above property model.

Table 1 Cap-Ex vs. Op-Ex Comparison (Infrastructure as a Service)

On Property	Yr1	Yr2	Yr3	Yr4	Yr5	Total
Licensing	8	8	8	8	8	39
Hosting	350	350	350	350	350	1,752
Support	312	322	331	341	352	1,659
Transaction Costs[1]	165	165	165	165	165	826
Upgrades	99	99	99	99	99	495
Depreciation	171	194	216	238	261	1,080
Total Annual Costs	**1,106**	**1,138**	**1,170**	**1,202**	**1,235**	**5,851**

Cloud	Yr1	Yr2	Yr3	Yr4	Yr5	Total
Licensing	0	0	0	0	0	0
Hosting	0	0	0	0	0	0
Support	619	619	619	619	619	3,097
Transaction Costs[1]	0	0	0	0	0	0
Upgrades	0	0	0	0	0	0
Depreciation	188	188	188	188	188	941
Total Annual Costs	**808** (-27%)	**808** (-29%)	**808** (-31%)	**808** (-33%)	**808** (-35%)	**4,038** (-31%)

Table 2 Example Infrastructure as-a-Service vs. On Property Cost Comparison

IAAS

	Yr1	Yr2	Yr3	Yr4	Yr5	Total
Network	-	$ -	$ -	$ -	$ -	$ -
Enterprise Licenses	$ 100,000	-	-	-	-	$ 100,000
POS Hardware	$ 1,072,985	-	$ -	$ -	$ -	$ 1,072,985
Implementation Services	$ 175,000	$ -	$ -	$ -	$ -	$ 175,000
Hosting	-	-	-	-	-	-
Server Hardware		-	-	-	-	-
Server	-					$ -
Total	**$ 1,347,985**	**$ -**	**$ -**	**$ -**	**$ -**	**$1,347,985**
Transaction Fees	-	-	-	-	-	-
SaaS Subscription	$ -	$ -	$ -	$ -	$ -	$ -
Annual Support	$ 16,850	$ 35,000	$ 35,000	$ 35,000	$ 35,000	$ 156,850
Hosting	$ 5,000	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 50,000
Support FTE & PCI Activities	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 250,000
Hardware Support	-	-	-	$ 15,000	$ 15,000	$ 30,000
Product Update Management	$ -	$ 8,000	$ 8,000	$ 8,000	$ 8,000	$ 32,000
Depreciation	$ 269,597	$ 269,597	$ 269,597	$ 269,597	$ 269,597	$ 1,347,985
Total	**$ 341,447**	**$ 412,597**	**$ 412,597**	**$ 427,597**	**$ 427,597**	**$1,866,835**

On Property

		Yr1	Yr2	Yr3	Yr4	Yr5	Total
Capital	Network	$ 75,978	$ -	$ -	$ -	$ -	$ 75,978
	Enterprise Licenses	$ 154,000	-	-	-	-	$ 154,000
	POS Hardware	$ 1,072,985	$ -	$ -	$ -	$ -	$ 1,072,985
	Implementation Services	$ 235,630	$ -	$ -	$ -	$ -	$ 235,630
	Hosting		-	-	-	-	$ -
	Server Hardware	$ 113,743					$ 113,743
Admin Capital	Server	$ 41,609					$ 41,609
	Total	**$ 1,693,945**	**$ -**	**$ -**	**$ -**	**$ -**	**$1,693,945**
Expense	Transaction Fees	-	-	-	-	-	-
	SaaS Subscription	$ -	$ -	$ -	$ -	$ -	$ -
	Annual Support	$ 16,850	$ 35,000	$ 35,000	$ 35,000	$ 35,000	$ 156,850
	Hosting	$ 75,000	$ 75,000	$ 75,000	$ 75,000	$ 75,000	$ 375,000
	Support FTE & PCI Activities	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 250,000
	Hardware Support	-	-	-	$ 15,000	$ 15,000	$ 30,000
	Product Update Management	$ -	$ 8,000	$ 8,000	$ 8,000	$ 8,000	$ 32,000
	Depreciation	$ 338,789	$ 338,789	$ 338,789	$ 338,789	$ 338,789	$ 1,693,945
	Total	**$ 480,638**	**$ 506,789**	**$ 506,789**	**$ 521,789**	**$ 521,789**	**$2,537,794**

Table 3 Example Dedicated Hosting vs. On-Property Costs

Dedicated Hosting						
	Yr1	Yr2	Yr3	Yr4	Yr5	Total
Network	$ 75,978	$ -	$ -	$ -	$ -	$ 75,978
Enterprise Licenses	$ 154,000	-	-	-	-	$ 154,000
POS Hardware	$ 1,072,985	-	$ -	$ -	$ -	$ 1,072,985
Implementation Services	$ 235,630	$ -	$ -	$ -	$ -	$ 235,630
Hosting	$ 110,000	-	-	-	-	-
Server Hardware	$ 113,743	-	-	-	-	-
Server	$ 41,609					$ 41,609
Total	**$ 1,803,945**	**$ -**	**$ -**	**$ -**	**$ -**	**$1,580,202**
Transaction Fees	-	-	-	-	-	-
SaaS Subscription	-	-	-	-	-	$ -
Annual Support	$ 16,850	$ 35,000	$ 35,000	$ 35,000	$ 35,000	$ 156,850
Hosting	$ 100,000	$ 100,000	$ 100,000	$ 100,000	$ 100,000	$ 500,000
Support FTE & PCI Activities	-	-	-	-	-	$ -
Hardware Support	-	-	-	$ 15,000	$ 15,000	$ 30,000
Product Update Management	$ -	$ -	$ -	$ -	$ -	$ -
Depreciation	$ 360,789	$ 360,789	$ 360,789	$ 360,789	$ 360,789	$ 1,803,945
Total	**$ 477,639**	**$ 495,789**	**$ 495,789**	**$ 510,789**	**$ 510,789**	**$2,490,795**

On Property							
		Yr1	Yr2	Yr3	Yr4	Yr5	Total
Capital	Network	$ 75,978	$ -	$ -	$ -	$ -	$ 75,978
	Enterprise Licenses	$ 154,000	-	-	-	-	$ 154,000
	POS Hardware	$ 1,072,985	$ -	$ -	$ -	$ -	$ 1,072,985
	Implementation Services	$ 235,630	$ -	$ -	$ -	$ -	$ 235,630
	Hosting		-	-	-	-	$ -
	Server Hardware	$ 113,743					$ 113,743
Admin Capital	Server	$ 41,609					$ 41,609
	Total	**$ 1,693,945**	**$ -**	**$ -**	**$ -**	**$ -**	**$1,693,945**
Expense	Transaction Fees	-	-	-	-	-	-
	SaaS Subscription	$ -	$ -	$ -	$ -	$ -	$ -
	Annual Support	$ 16,850	$ 35,000	$ 35,000	$ 35,000	$ 35,000	$ 156,850
	Hosting	$ 75,000	$ 75,000	$ 75,000	$ 75,000	$ 75,000	$ 375,000
	Support FTE & PCI Activities	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 250,000
	Hardware Support	-	-	-	$ 15,000	$ 15,000	$ 30,000
	Product Update Management	$ -	$ 8,000	$ 8,000	$ 8,000	$ 8,000	$ 32,000
	Depreciation	$ 338,789	$ 338,789	$ 338,789	$ 338,789	$ 338,789	$ 1,693,945
	Total	**$ 480,638**	**$ 506,789**	**$ 506,789**	**$ 521,789**	**$ 521,789**	**$2,537,794**

Table 4 Example Software as-a-Service vs. On-Property Costs

SAAS						
	Yr1	Yr2	Yr3	Yr4	Yr5	Total
Network	-	$ -	$ -	$ -	$ -	$ -
Enterprise Licenses	-	-	-	-	-	$ -
POS Hardware	$ 954,720	-	$ -	$ -	$ -	$ 954,720
Implementation Services	$ 100,000	$ -	$ -	$ -	$ -	$ 100,000
Hosting	-	-	-	-	-	-
Server Hardware		-	-	-	-	-
Server	$ 41,609					$ 41,609
Total	**$ 1,096,329**	**$ -**	**$ -**	**$ -**	**$ -**	**$1,096,329**
Transaction Fees	-	-	-	-	-	-
SaaS Subscription	$ 150,000	$ 150,000	$ 150,000	$ 150,000	$ 150,000	$ 750,000
Annual Support	-	-				
Hosting	-	-	-	-	-	-
Support FTE & PCI Activities	-	-	-	-	-	$ -
Hardware Support	-	-	-	$ 15,000	$ 15,000	$ -
Product Update Management	$ -	$ -	$ -	$ -	$ -	$ -
Depreciation	$ 219,266	$ 219,266	$ 219,266	$ 219,266	$ 219,266	$ 1,096,329
Total	**$ 369,266**	**$ 369,266**	**$ 369,266**	**$ 384,266**	**$ 384,266**	**$1,846,329**

On Property							
		Yr1	Yr2	Yr3	Yr4	Yr5	Total
Capital	Network	$ 75,978	$ -	$ -	$ -	$ -	$ 75,978
	Enterprise Licenses	$ 154,000	-	-	-	-	$ 154,000
	POS Hardware	$ 1,072,985	$ -	$ -	$ -	$ -	$ 1,072,985
	Implementation Services	$ 235,630	$ -	$ -	$ -	$ -	$ 235,630
	Hosting		-	-	-	-	$ -
	Server Hardware	$ 113,743					$ 113,743
Admin Capital	Server	$ 41,609					$ 41,609
	Total	**$ 1,693,945**	**$ -**	**$ -**	**$ -**	**$ -**	**$1,693,945**
Expense	Transaction Fees	-	-	-	-	-	-
	SaaS Subscription	$ -	$ -	$ -	$ -	$ -	$ -
	Annual Support	$ 16,850	$ 35,000	$ 35,000	$ 35,000	$ 35,000	$ 156,850
	Hosting	$ 75,000	$ 75,000	$ 75,000	$ 75,000	$ 75,000	$ 375,000
	Support FTE & PCI Activities	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 50,000	$ 250,000
	Hardware Support	-	-	-	$ 15,000	$ 15,000	$ 30,000
	Product Update Management	$ -	$ 8,000	$ 8,000	$ 8,000	$ 8,000	$ 32,000
	Depreciation	$ 338,789	$ 338,789	$ 338,789	$ 338,789	$ 338,789	$ 1,693,945
	Total	**$ 480,638**	**$ 506,789**	**$ 506,789**	**$ 521,789**	**$ 521,789**	**$2,537,794**

3.1 Dedicated Hosting

Dedicated hosting as compared to an on-property solution is a route many companies choose to take in order to reduce their physical footprint or shift roles and responsibilities for an IT organization.

As a company, you want to look at the financial differences between hosting in your data center and having a third party host your POS server infrastructure. You cannot omit your physical POS footprint in this scenario; only server infrastructure. In both scenarios, you have network, licensing, hardware, services, hosting and support costs. The only area where there may be savings is on the side of support and responsibility. Most times, a third party engagement can assume the support responsibility for the infrastructure/hosting center only. The application layer is still the responsibility of the hotel.

Lastly, the third party provider could assume any PCI or other regulatory services from a hosting perspective. Each element needs to be contractually defined. However, there may be additional implicit cost savings for an organization.

Overall, it is not typically seen as a direct cost savings approach to simply move to dedicated hosting. The decisions to use dedicated hosting are mostly around other business elements not driven by finances. There is generally no major change in capital vs. expense with this model.

3.2 Infrastructure as-a-Service (IaaS)

Infrastructure as a Service as compared to an on-property solution is a route many companies choose to take in order to reduce their physical footprint and save money for an IT organization.

You want to look at the financial differences between hosting in your data center and using infrastructure as a service to host your server environment. The major difference financially with IaaS providers is the mentality of “pay for what you use.” The idea behind this for a POS infrastructure is to be able to “turn down” or lighten infrastructure use during non-peak hours to reduce costs. This is the primary focus for direct cost savings with this model. The secondary aspect is the reduced need to purchase physical hardware. The reason it is secondary is because it is a reduction in capital spend. The primary cost savings driver is an expense model. Looking at these costs over a five-year period could reduce your costs significantly even when your operating model is more expense driven.

An organization leveraging IaaS has the implicit responsibility for the environment within. Unlike dedicated hosting or SaaS, the hotel/organization is generally responsible 100% at the application and network layer. There are generally no contractual considerations beyond the infrastructure layer. That said, the hosting aspects to PCI (as an example) are generally significant which may make IaaS a feasible option.

Overall, there are some responsibility definitions and understanding with IaaS. Your financial model will become expense heavy. For a project implementation/justification, you need to expand beyond one year to see the justified spend and/or savings.

3.3 Software as-a-Service (SaaS)

Software as a Service as compared to an on- property solution is a route many companies choose to take in order to reduce physical footprint, shift roles and responsibilities, and save money for an IT organization.

As a company, you want to look at the financial differences between hosting in you data center and using SaaS to manage your environment. The key word with SaaS is "manage," unlike with the previous two models. This model is typically seen as a 100% shift of management and responsibility from the power to service delivery. A company generally does not concern themselves with the workings of a SaaS provider given they have appropriate certifications and the appropriate contractual agreements are in place.

Financially, SaaS is heavy on the expense side. Like IaaS, for project implementation/ justification, you need to expand beyond one year to see the justified spend and/or savings.

In short, the more you are responsible for the less expensive a provider can be. However, there is a cost to taking on more responsibility. At the end of the day, as an organization you are ultimately responsible for all elements of a system's infrastructure and where it is hosted. Finances are only part of an evaluation and should always be considered with the technology and operations of said systems. Up front expenses will be higher than with a traditional on-property model.

Figure 1 Example Operational Comparisons

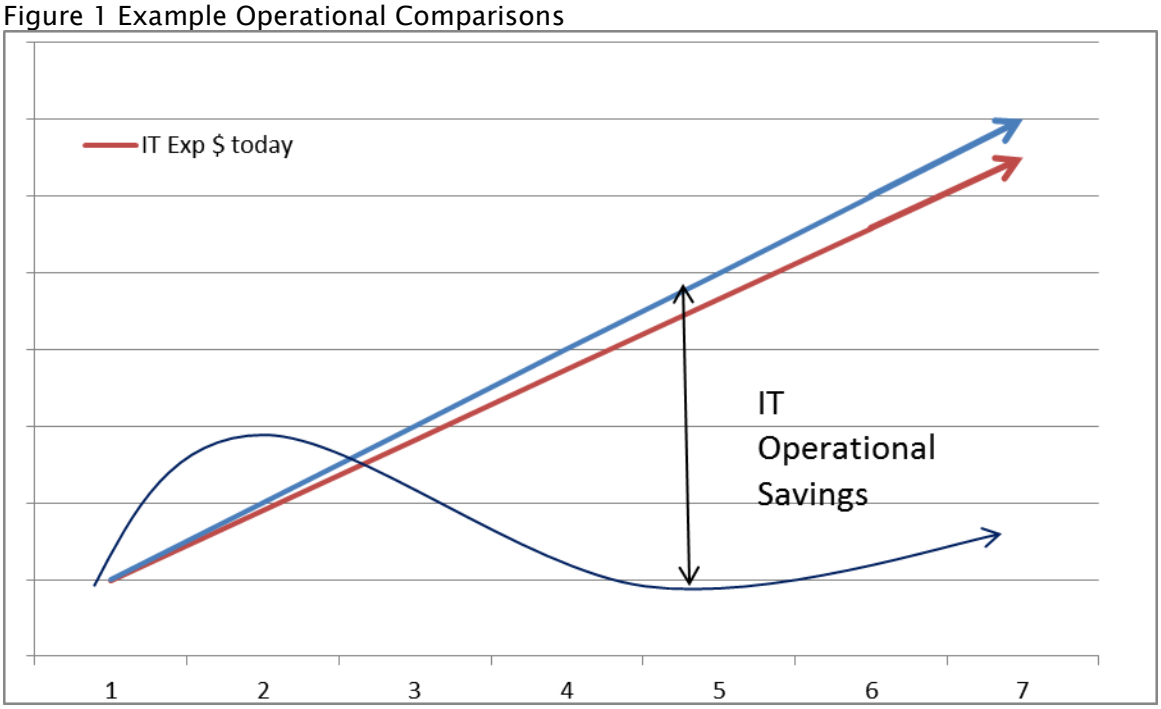

3.4 Summation

Sometimes the numbers do not work out to justify cloud. However, there's a value-based cost which needs to be considered. If an organization no longer needs to be concerned with system patches/maintenance, hardware maintenance, upgrades and general day-to-day support, what is that worth to the organization?

What is the value to you as an organization:

- To not have to worry about supporting a hosting location?
- To not have to worry about upgrading your environment with support staff?
- To not have to worry about support of an enterprise system in the cloud?

There is inherent value with utilizing cloud services. This can often be placed into your TCO model as a reduction. However, it is found there is never really a reduction. Rather, there is a shift in focus for your head count. Your staff is more readily available to act strategically and be innovative. Your staff is available to build new functionality into your cloud services. They are no longer constrained with needing to focus on day-to-day support of said systems. This shift in focus has many benefits with SaaS type solutions:

- Focus on delivering IT services "at the speed of business"
- Transform into leaders for new technology
- Foster new ideas and new ways of doing business
- And of course, provide a great experience to your customers.

Of course, adoption of new technological methodologies will never occur unless you have buy-in and understanding from both your finance and operational staff. Easier said, if you cannot justify the costs and the system(s) does not function to meet the needs of your business, the technology will be a non-starter. This serves as a checks-and-balance for technological requests as well. New technology is great. However, if it's not going to operationally function there's no use for it. If the costs are too askew, you may not want to afford it. The same is true with any above property evaluation. Does the technology exist? Does it financially make sense? Can it operate? Then you may have the beginning of something you can use.

4 Hosting Considerations

While this section will by no means provide you with a comprehensive listing of hosting considerations, there are five key areas that warrant special attention during the evaluation process of whether or not a hosted application environment will make sense for your organization.

4.1 Facilities

The first question to ask is "who?" Who will provide the hosting facility? Will the physical location that will house the application computing environment (the data center) be a facility managed by your company, by the application vendor or by a third party who may specialize in data center operations?

If your organization already maintains a robust data center environment, complete with knowledgeable and trained technology resources, then adding an additional application to the mix may not be an undue burden on your organization. From an economies of scale perspective, this option may make sense if the foundational data center requirements are already being met. In this model, you likely already have established robust and secure data connectivity between your properties and your data center so the requirement of required bandwidth and network speed may already be met.

If the application vendor partner is offering hosted services, this may provide an attractive option for you as well. In this model, you are placing all of the responsibility, and trust, in your vendor partner to deliver not only a functioning application, but also a professionally-managed data center environment. Many refer to this model as the "one throat to choke" approach where, when issues arise, there is but one call to place to report your issue.

A third approach provides a hybrid solution between the first two. Many international computing leaders are now offering third-party data center facilities and services for either your organization, or for your vendor partner, of which to take advantage. As these organizations, such as Amazon and AT&T, do this for at scale, oftentimes they can offer attractive hosting service pricing as a means of lowering your total cost of ownership. In this model, however, there are three parties involved: the property, the application vendor partner and the data center provider. Support and troubleshooting of reported issues can become complicated quickly if extreme care is not taken up front to define roles and responsibilities of all parties.

Regardless of which model you choose, it is important to assess what your organization's true core competency is and then select a model that plays best to your strengths. If your organization is a hospitality organization, then establishing a data center operation to host applications may not play to your strengths. If you are a multi-national organization, with a variety of enterprise-wide applications that you host in data centers across the globe, then adding one more application to the mix may be something to seriously consider. If price is a

motivating factor, and the potential for support complexities is not a large factor, then the hybrid third-party data center solution may be most appealing.

4.2 Location

Where the data center facility physically resides is the next consideration. If your operation is concentrated in a small geography, say perhaps the southern half of France, then identifying a hosting facility within that same geography is a logical choice. If those same properties in France are being asked to connect to a data center in Mexico or in Singapore, then you can plan to expect greater latency between the time your users hit the "enter" key and the application returns a response. Not all hosted applications have the requirement today to send each transaction to the data center and back before allowing your users to continue on with their work, but many applications still do. You will want to check on this.

If your estate of locations is geographically diverse, does it make sense to have all of your properties report back to a single data center? Building upon the example above, if all of your locations are based in the European Union, then perhaps a data center in Germany is a good solution to consider.

If your estate of locations is scattered across 75 countries around the world, then perhaps a regional data center approach may be more appealing. Employing this model, data will not need to travel as far to reach a single data center. It may be logical to group all of your regional properties together in a regional data center to provide you with regional reporting solutions. With this approach, however, special plans will need to be made to bring all of your regional data together to produce true enterprise reporting across all of your regions.

A final, yet the most important, consideration as to where a data center is to be located involves your need to adhere to country-specific privacy laws. A number of countries today mandate the do's and don'ts of what data can cross borders on its way to and from data centers. Special attention will need to be paid to ensure that you are aware of the privacy laws in place within each geographic location where your company conducts business.

4.3 Support

As mentioned above, great care should be taken to thoroughly define the roles and responsibilities of all parties engaged in the design of the hosted landscape. The more parties that are involved in providing a piece of the overall solution, they greater the complexity will be in defining the support model. What is key here is to have a full understanding of not only who is responsible for what, but also who is responsible for the overall management of the support process.

Many hosting agreements will place support responsibilities on the vendor partner only after the data message has been received at the vendor's data center facility. The local area network at your facility, the establishment of Internet connectivity and the physical act of getting the data message from the property to the vendor's data center is often the property's

responsibility. When the Internet connection is lost, when a POS receipt printer gets dropped into a vat of grease or when none of your computer workstations at the hotel will talk with one another, who will you call for support? While the application vendor partner may provide support services for all of these components, it should be verified and contractually defined up front. In many cases, additional support contract agreements may need to be entered into for Internet or hardware support services. Regardless of the mix of support providers, ensure you have a comprehensive look at your total cost of support services.

4.4 Connectivity and Data Transmission

The transmission of data from the application workstation to the hosting center will often become a property responsibility. While application providers will often publish minimum connectivity requirements in terms of bandwidth or message latency metrics, a savvy consumer will want to check with references to ensure these stated “minimums” will really address the business needs and render the application functional from a practical standpoint. If your users are waiting 10-20 seconds for the application to return data with each and every transaction, then bandwidth needs are not meeting your operational requirements.

A secondary, or failover, Internet connection may also be considered if the application is deemed mission critical in nature. Should your primary Internet connection at your property fail, will you want a mechanism in place to automatically switch over to a fallback Internet connection? In many locations, the cost of bringing a second circuit from an alternate Internet provider is a prudent insurance policy for your operation. Your guests and employees both will thank you for planning ahead.

So connectivity has been established, a secondary Internet connection stands at the ready, and you have confirmed your bandwidth and latency metrics exceed minimums. All is well on the connectivity front, at least for a while. When connectivity issues do present themselves, (and they eventually will) how will you address sudden slowdowns or other latency issues? Until the data message physically arrives at the hosting center, your data center partner will not be in a position to assist you with the troubleshooting steps necessary. Where is your data going? How many segments, or hops, is the data making between switches and other facilities between the time it leaves your property workstation and when it arrives at the data center? Will your Internet service provider assist you in diagnostics and troubleshooting? Has this been accounted for in your service contract with them?

How can you be certain that any latency is not occurring within the data center facility? It is entirely possible for the data to get from your property to the data center within acceptable latency guidelines only to be bogged down by inefficient data processing practices. Does your data center have the ability to proactively monitor their data center network and server resources? Is there contention between the database server and the application server? Do only certain types of application transactions slow the entire system down? Great levels of complexity will likely greet you in this area. Having a plan up front, and understanding which of

your partners are equipped with the correct diagnostics tools to assist in the troubleshooting, will alleviate chronic issues in this area.

A final thought in this area: work with your Internet service provider to understand where your data is going on its way to the data center. In a recent real-life scenario, message latency was being experienced at a property. Their hotel was located on the East Coast of the United States and the data center itself was also located on the East Coast. Both locations were in major metropolitan areas with robust infrastructure.

Upon weeks of scrutiny, it was determined that the message was leaving the hotel, going to the Internet service provider's local facility, passing on to a second facility on the West Coast of the United States and then shooting back across the country to the application vendor's data center. What should have been at best a 200 mile journey was averaging over 6000 miles with each message sent.

4.5 Disaster Recovery

This topic alone will get the attention of your business partners when you first approach them with the idea of considering a hosted application. Visions of a busy Saturday night at your facility – with guests everywhere and a booming business in every corner of the property – and a computer system that is not working due to lack of connectivity to a data center can be downright paralyzing. Having a proper and comprehensive disaster recovery plan in place will be the answer to assuage any operational fears. Three distinct areas of disaster recovery should be considered.

What happens within your property is paramount to your operational wellbeing. If data center connectivity is lost, will your application continue to function? Many of the newest applications are being architected in such a way that data center connectivity is no longer a requirement for the application to function as designed. Transactional data is retained at the workstation level and then passed along to the data center when connectivity is restored. Other legacy applications may not be architected in this fashion; once data center connectivity is lost the application is rendered useless.

A second strategy to consider is to have a standby connectivity option available. Some locations that are prone to natural disasters, for example in an area prone to hurricanes, terrestrial Internet connectivity may be lost for multiple days. By attaching a 3G/4G cellular adapter to a workstation, it may be possible to re-establish connectivity to the data center utilizing this alternate approach. While it may not bring all workstations back online, it will, at a minimum, provide the operation with a means to continue processing the business transactional needs. It will be important for you to discuss what options are available to you with your vendor partner. If their legacy versions of application cannot provide you with an on-property disaster recovery solution, then perhaps their latest generation of product offerings will. An application upgrade may be all that it takes to give you expanded disaster recovery options.

The second area of disaster recovery involves getting the data between your property and the data center. As discussed earlier, having more than one Internet connection in and out of your property serves your disaster recovery needs well. Investing in network technology that will automatically connect from the primary to the secondary circuit in the case of a disaster will save you operational downtime.

An additional component to investigate involves an understanding of what disaster recovery measures the data center facility has taken within their walls. Do they have redundant switches, firewall, routers and load balancers? If one of those units goes down, does your application come to a halt? What about the application and database servers themselves? Will your properties be connecting to a single server solution or are there failover, or cluster server hardware resources, built into the overall data center design. Ensure these details are negotiated and defined in any hosting agreement.

When the unthinkable happens and there is a catastrophic failure at the data center facility, what sort of failover disaster recovery is in place to get you up and running again at a secondary data center facility? Does your hosting agreement make provisions for real-time or near real-time data backups between your primary and a secondary data center facility? Will similar hardware be provided at this secondary data center facility or will your properties be asked to run on a scaled down infrastructure? Will you have access to 100% of your data and application functionality? 80%? 50%? It is very possible in today's environment to ensure secondary data center facilities are equipped and standing by as a third tier of disaster recovery planning. While all of this is possible, it will come with a cost. Factor the cost of a three tier disaster recovery solution against the criticality of the data served by this application. If the application is deemed "mission critical," then logic dictates that data has a higher value to your organization and should be protected with a higher level of disaster recovery plan. Weighing risk against cost will enable your organization to determine the proper mix of disaster recovery preparations.

Having a disaster recovery plan defined and in place can be seen as a major accomplishment in this journey. It will, however, not do you much good if it does not work.

Ensure in your negotiations with the hosting solution provider that a periodic disaster recovery drill is executed. Many organizations elect to conduct such drills on a semi-annual or annual basis. The frequency will be driven by how critical the application is seen to your organization. During a disaster recovery drill exercise, all aspects of your plan should be exercised. If you have elected to purchase 3G/4G cellular adapters to provide alternative connectivity to the data center, then take your Internet service offline at the hotel and watch as your operators source the adapter, install it and restore connectivity. Did they know what to do and where to find the adapters? Does on-site documentation provide clear instructions even your newest employee can follow?

If you employed network technology to automatically fail over your primary Internet connection to its failover equivalent, did that solution work when the primary Internet connection was rendered inoperative?

Working with the data center and/or application provider, plan an exercise where the application database server is taken offline. Did a failover server kick in? Did any proactive monitoring alerts get triggered to inform the data center team of an issue? Were support tickets opened as expected? Did any email escalations initiate as designed?
Finally, if your application is setup to failover to a secondary data center, this exercise must also be practiced and validated. Once the secondary data center has become the primary data processing location for your properties, can you fall back to the original primary data center? What was the timing of operational downtime between when your property was connected to the secondary (failover) data center and when it was connected back to the original primary data center. Planned failover times should be yet one more metric to include in your disaster recovery agreement.

Ultimately, there are several aspects to take into consideration for quick recovery in the event of disasters. As more software moves into the cloud, customer failures begin to overlap with the provider's. SaaS helps to solve more of the business-side disasters, but has created more problems when a client experiences failures.

One example exists within the VPN/ACL as user IPs can be dynamic and frequently change. The solution here is to attempt to utilize a Domain Name System (DNS) in order to assign domain names to any resource connected to the VPN. This helps to organize and identify the many devices that may be connected to the network. Another solution is to pre-plan for any failures in order to efficiently respond to user problems. With foresight, hotels can build their data centers into the access control list (ACL) in order to have all rules and access in place for failover locations.

An alternative recommendation is to conduct disaster recovery drills to always be ready in the event of actual problems. More commonly these days network access is required for many key systems, and it is important to plan to have redundant ISP connections in case of failure at one. A very basic preventative measure is to ensure that the core network infrastructure is on Automatic Protection Switching (APS) and generator power.

5 Other Considerations

5.1 Testing

Similar to training employees on the software, it is imperative to also run the product through a series of tests in order to ensure it is operational. There are three stages of software testing to run through in order to validate its effectiveness: Beta, Functional and Integration testing.

Beta testing in what is called a “sandbox” can be very helpful in rooting out bugs and glitches. This is a way to test features before they go live, which can be more of a challenge with SaaS. When implementing functionality that is new for customers, it is important to have a migration plan in place for any existing customer data affected by it. Likewise, there should be a rollback plan to minimize impact on customer data. Many providers perform a production deployment with switches to enable a new feature selectively for testing.

Functional testing allows providers to examine whether the product does what it is meant to do. This process offers a space to discover whether there will be an impact on existing user workflow and determines whether the user will need to be trained on or before the rollout. Once again, it is helpful to have a training area where new functionality can be tested without impacting data in production.

Integration testing gives insight to how the product will work with various external elements. A basic checklist for integration testing is whether the product works with other systems (i.e. APIs), other OS and/or browsers, and the total working environment. This gives the developers a platform to observe how the software performs against peripheral factors.
These three methods of testing are most useful when using live data. This allows developers to “load-test” and see exactly how the software responds to the messiness of “real” data and actual load of real usage. Paired with the training methods mentioned above, the software can be tested against nearly every scenario that may befall it in reality.

5.2 Monitoring and Availability

After development and testing is complete, a significant amount of monitoring and availability should be put into the final product. When monitoring the software, a periodic check of network availability and performance, as well as a way to track changes through an audit trail, is necessary. Availability should be examined through the lens of infrastructure, system and data center, or geographic redundancy.

To monitor network availability one should have a real-time scan of the system’s uptime. If there is downtime, it should be reported instantly for support and diagnosis. Performance monitoring should help expedite troubleshooting, reduce downtime and have alerts configured to notify when a problem arises. An audit trail can help troubleshooters look back at the history of the network to further isolate the origination of problem(s).

One should be consistently aware of how available the network is and how each step within the network is performing. By having indicators to signal even the smallest errors, larger problems can be caught before they begin. By incorporating tiers into data centers, it is easier to track uptime and network resiliency. Typically, networks are set up to have 1-4 tiers to hierarchically grade the network. An example of this tiered process follows[2]:

Table 5 Backup and Availability

	Backup Method	Guaranteed Availability
Tier 1	Single uplink & servers	99.671%
Tier 2	*Tier 1* + Redundant capacity components	99.741%
Tier 3	*Tier 1* + *Tier 2* + Dual powered equipment and multiple uplinks	99.982%
Tier 4	*Tier 1* + *Tier 2* + *Tier 3* + all components are fully fault-tolerant including uplinks, storage, chillers, HVAC, servers, etc. and is dual powered	99.995%

It is important to analyze even the most basic pieces of data center infrastructure including making sure it is on multiple power grids with multiple ISP connections, identified as Tier 4 in the table above. Even having the utilities such as HVAC in proper operating conditions and redundant is critical. At the system level again having multiple power supplies is important, as well as the inclusion of multiple hard disks (RAID or SAN) and multiple servers if possible. And if the network spans multiple geographies it is critical to duplicate their setups and maintenance so if one fails the others can be used as backups.

[2] Craft, Nix. "Explain: Tier 1/2/3/4 Data Center." *Linux Tutorial for Beginners and Advanced Users, NixCraft RSS*. Cyberciti. June 7, 2008. Web. 23 Sept. 2014. <http://www.cyberciti.biz/faq/data-center-standard-overview/>.

5.3 Public and Private Web Services

Public and private web services necessitate different security structures. Publicly accessible web services that are open to anyone on the web require a significantly greater amount of security and must protect themselves from denial-of-service (DoS) attacks. Although the consequences of DoS can vary, there is a high risk of interrupted host services.

One technique often used to keep public services more protected is a secure token or credential for access to service. This authentication can be provided on a temporary basis and with specific, limited access. Another useful method is to apply metering, or what is sometimes referred to as governors. These limitations set on the service track and enforce certain statistics to regulate usage. Salesforce.com, for example, outlines some of the statistics that can be tracked, including: per-transaction limits, size-specific limits, and also email and push notification limits. [3]

Private web services are much more secure and offer a more robust barrier to malicious attacks. Private services such as virtual private networks (VPNs) require a level of authentication to access and can only be reached via a secure network. VPNs are commonly used for business intranet that allows employees to communicate remotely and safely across a large network. Although private services deliver a higher level of protection to a host, they require additional infrastructure that is associated with higher costs.

5.4 VPN Models

Two popular models of enterprise VPNs exist that each possess their own individual strengths and weaknesses. The hub-and-spoke model, sometimes referred to as one-to-many, occurs when one data center is connected to several individual sites. This model is relatively dependable, as one failure does not bring the entire system down. By having the individual sites disparate and spread out, the hub-and-spoke model offers a form of redundancy in the case of a failure. One drawback of the hub-and-spoke model is that it can be fairly difficult to spot specific few outages, as the network can be far-reaching and widespread.

An alternative VPN topology is the LAN-to-LAN, or hub-to-hub, model. This model is more reliable with one entire data center communicating with another. One hub will have to be identified as the VPN server while the others will connect to it. A benefit to this setup is that hub-to-hub VPNs allow isolation of a single point of failure. Since many sites within the

[3] "Understanding Execution Governors and Limits." *Understanding Execution Governors and Limits*. Salesforce, n.d. Web. 19 Sept. 2014.
<https://www.salesforce.com/us/developer/docs/apexcode/Content/apex_gov_limits.htm>.

particular LAN would be affected in the event of failure, it is easy to determine which hub is causing problems. [4]

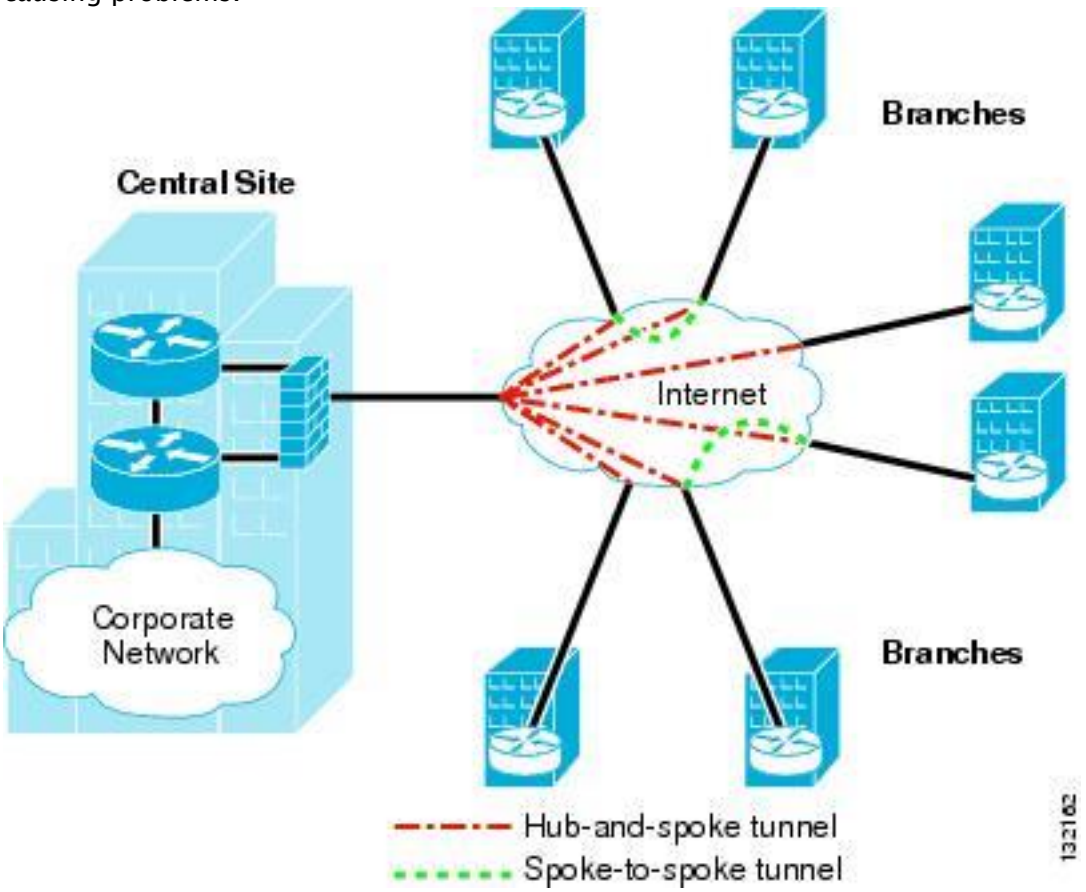

Figure 2 Example VPN Models (source: www.cisco.com)

5.5 Latency

Latency is the time delay that results from data being transmitted across the physical points in a network. Latency can be measured by administering ping tests or a traceroute, a diagnostic tool that allows one to measure the path of the packets being sent across the network. Contributors to high latency are transmission and propagation delays of data, as well as the number of routers and other hardware delays. These various aspects can all lead to more network congestion resulting in higher latency issues. [5]

[4] "10.1 Types of VPNs." *PacketiX VPN 2.0 Manual*. Plat' Home. Web. 22 Sept. 2014 <http://www.plathome.com/support/packetix/manual/10-1.htm>.

[5] "Network latency effects on application performance." *SearchITChannel*. TechTarget, Nov. 2006. Web. 18 Sept. 2014. <http://searchitchannel.techtarget.com/feature/Network-latency-effects-on-application-performance>.

Typically, these contributors rise in networks and applications that are geographically and/or geospatially separate, forcing data to travel over a large physical distance. Some examples are satellite connections and international customers far from the originating data center. In both of these settings, data is travelling for a relatively long period of time causing the end user to experience slower connection speeds. To prevent this frustration, a network host can implement a content delivery network (CDN) or caching content to a local disk, covered in the following two topics.

5.6 Caching

Caching allows frequently accessed information to be disseminated across networks. This is especially helpful for both large and small resources (files, databases, etc.) that are frequently needed. Caching can be done at the user level by web browsers, at the LAN level by dedicated caching appliances and software, and at the WAN level by similar appliances and software in distributed data centers. There are many approaches and architectural models to caching that can be done at the network, application or resource level. The overall benefit of caching is that it allows faster access to resources by minimizing the network path the resource needs to travel to get to the user. The caching and cache expiration strategies employed must align with the expected use and timeliness of the resources being cached.

A simple example of local (browser-based) caching follows:

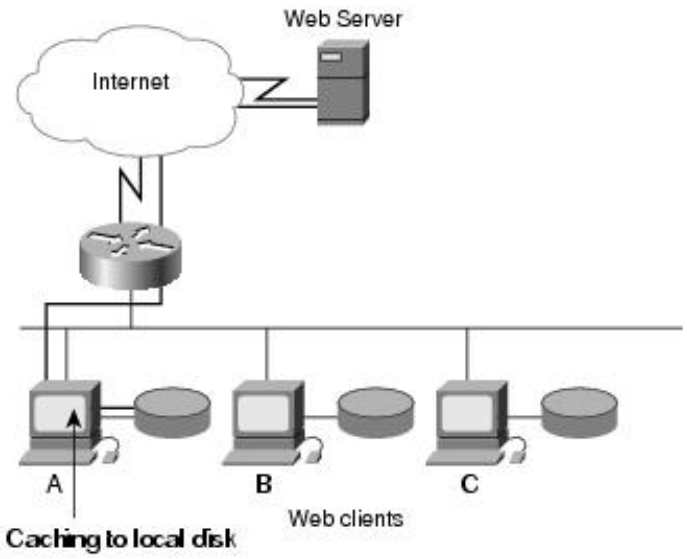

Figure 3 Browser-Based Caching (source: www.cisco.com)

5.7 Content Delivery Networks (CDN)

Content delivery networks (CDNs) are a caching strategy used by web publishers and web applications to allow a host to offer content closer to where it is needed. Typically deployed in a series of nodes over multiple locations, CDNs offer a variety of benefits including bandwidth cost reduction, decrease in load times and a global availability of content. Content is cached in various locations across the globe, and they allow fast and reliable retrieval when necessary.

Larger CDNs make it possible for a provider to transfer its content to many devices efficiently, and alleviate problems that arise from limited bandwidth or demand spikes. Similar to caching, this occurs as the CDN delivers content stored on a server closest in proximity to the user. The process is essentially transparent to the user and happens almost instantaneously.

There are two types of CDNs, each with their own strengths and weaknesses. A push CDN allows a publisher to upload content to the network and link to and from it. The benefit with a push CDN is that the publisher chooses what content is uploaded and when it is updated or removed. This method also requires a minimum amount of bandwidth from the distribution center as this traffic is only generated when publisher choose to publish the content.

The second type of CDN is a pull CDN, where content is maintained on the distribution server and linked from the CDN by the publisher. Once retrieved, the content is cached by CDN. Pull CDNs are more relevant for web applications and other highly dynamic or custom websites that have a large amount of static data (training videos, etc.). [6] A single website or application can employ both push and pull CDNs in its architecture.

[6] Claire. "How to Choose the Right CDN For Your Website at WhoIsHostingThis.com." *Who Is Hosting This: The Blog*. Who Is Hosting This?, 30 June 2010. Web. 23 Sept. 2014. <http://www.whoishostingthis.com/blog/2010/06/30/cdns-push-vs-pull/#.>.

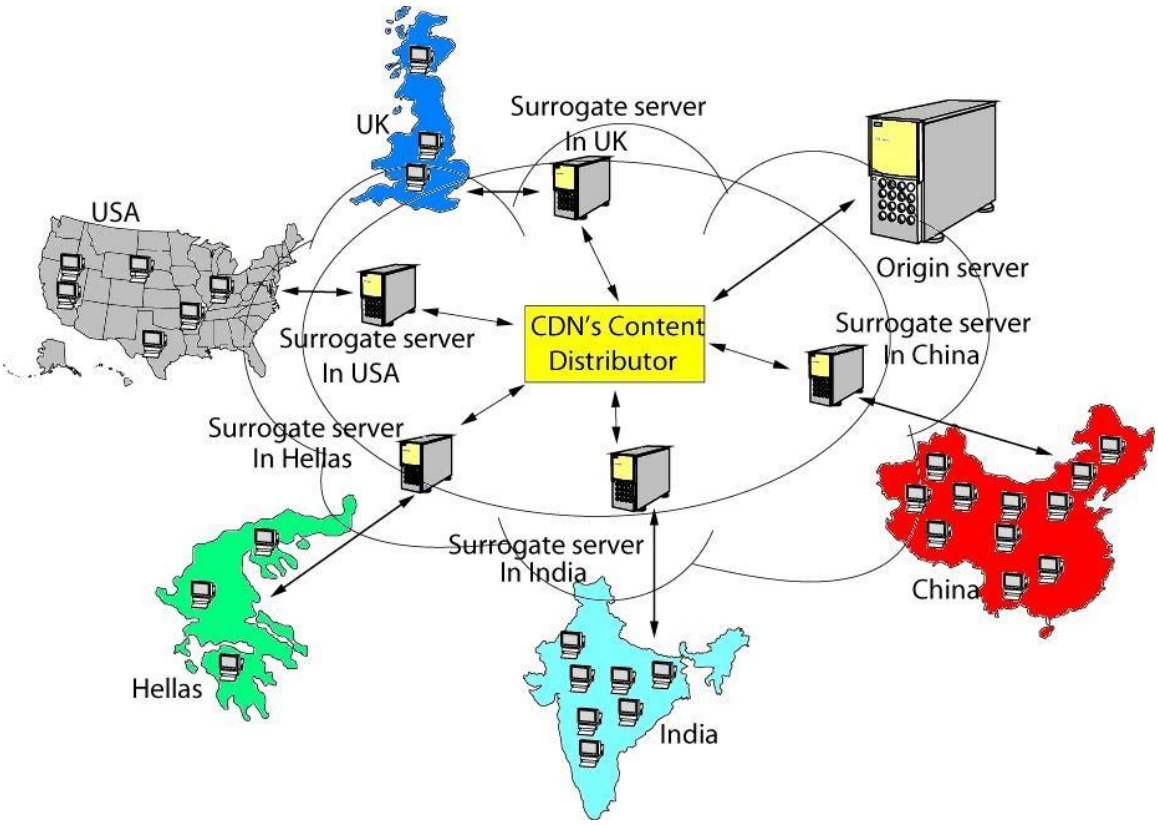

Figure 4 Example CDN (source: http://amkpathan.wordpress.com/article/ongoing-trends-and-future-directions-in-3uxfz2buz8z1w-2/)

5.8 Ports and Protocols

At the base level, there are two main types of Internet protocol traffic: Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). These core protocols act as managers of the information sent by a user. There are several differences between the two protocols, mostly revolving around reliability of the transfer. TCP is considered more reliable as long as a connection exists. No corruption occurs while transferring data, whereas with UDP corruption may occur. UDP also does not guarantee delivery of the data. [7]

Ports exist as communication endpoints in a host's operating system. Literally, a port is a 16-bit unsigned integer ranging from 0 to 65535. Every process over a network has a port to fall into. Well-known, or system, ports are listed between 0 to 1023, including HTTP, HTTPS, FTP and SSH. Registered ports range from 1024 to 49151 and are assigned by Internet Assigned Numbers Authority (IANA) for a specific service, commonly used by specific applications such as

[7] Craft, Nix. "What Is the Difference between UDP and TCP Internet Protocols?" *Linux Unix Tutorial for Beginners and Advanced Users NixCraft RSS*. Cyberciti, 15 May 2007. Web. 23 Sept. 2014. <http://www.cyberciti.biz/faq/key-differences-between-tcp-and-udp-protocols/>.

CDNs, management consoles and multiplayer games. Ports 49152 to 65535 are dynamic, private or ephemeral ports that cannot be registered with IANA.

It is important to understand the ports and protocols your application use in order to properly configure your security systems to allow this traffic. Also, protocols differ in their ability to handle network latency and interruptions. Knowing what type of traffic your above-property applications are generating allows you to properly build and configure the network to support it.

5.9 Identity Management

Identity Management is an enterprise's strategy to control and track its user's systems access. Due to high employee fluctuation and low resources, hoteliers have a special interest in high security and efficiency schemes to run systems and solutions. You need to know who or what is accessing your environments and actively enforce employee sign-on to your HR application or passively log visitor activity on your public website. Your focus with Identity Management is with authentication, which is what happens when people or machines log in and are identified as permitted to access.

5.9.1 Target

Security is your most important objective. You need to protect your assets, such as guests, employees, data, revenue streams, knowledge, infrastructure, etc. Efficiency is your key to high availability, lean management, automation and overall satisfaction. These targets can only be met by solutions providing customer user account directory based authentication technology.

5.9.2 Authentication Methods

In the ideal world, we have one single user account directory and all solutions accesses are based on it. We can deactivate an account and all access will be denied and the user has the luxury of a single-sign on. It is important to differentiate between enforced log-in where a user needs to type in credentials or where computer stored credentials are used to automatically log in. Where the enterprise may be at risk, the first way is preferred.

Sign-on processes and protocols:

- Federated authentication using SAML (Security Assertion Markup Language) allows you to send authentication and authorization data between affiliated but unrelated web services, enabling sign-on to the solution from a client application.
- Delegated authentication single sign-on enables you to integrate the solution with an authentication method that you choose, such as LDAP (Lightweight Directory Access Protocol) server or SAML, or perform single sign-on by authenticating using a token instead of a password. Delegated authentication provides stronger security and is preferred. SAML provides more flexibility for the solution provider.

5.9.3 Network

Your user account directory server needs to be behind a firewall that is open to communicate with the solution and vice versa. If your enterprise is at risk, the solution visibility may additionally be restricted to your network (see network chapter, IP and proxy restriction etc.).

5.9.4 Log-in provider

With today's solutions and applications, mainly those having browser-based front ends should be selected. With this, most computers and mobile devices can be used, and it is the most common standard. Of course, there are many other means, such as meta-frame solutions like Citrix or Terminal Server providing secure access, but these should only be applied where browser is not possible. In order for a user to log in to a web page, fields to type in credentials are required. This is the log-in provider. If you wish to allow credentials to be saved in the browser with the particular log-in provider is your decision. As an alternative, the browser's own log-in solution may be used, but this is less controllable.

5.9.5 Topology

The key question is if you have the log-in provider in your environment or with the solution vendor's.

It may be preferred to have it with the vendor as it will save you otherwise necessary web server infrastructure.

5.9.6 Directory and Solution Structure

Avoid any manual management of accounts or attributes inside the solution in order to have everything automated on the solution side. If it is necessary to add attributes or information from third resources, automate it with a structured file extract and import procedure.

The solution and its application will have different basic user roles. In a PMS (Property Management System) for example, you have Front Office, Reservations, Sales etc. These basic roles have to be mapped with your user account directory, typically with a DG (Distribution Group) for a role. The permission level and sub roles will still be configured inside the solution.

If you operate several hotels, they need to be separated in OU (Organization Units) and identified with a unique code as an attribute.

5.9.7 Responsibility and Support

A consequence of providing your user account directory is that you are responsible for account management and high availability. You will have to provide a support organization serving as the first contact for request and help calls.

5.10 Critical Service Levels

The following service levels and explanations are common in the industry, but the specifics and details of service levels should be defined in contracts and other documentation.

"Severity 1" means an Incident which is business impacting or poses an imminent impact; full hosting center outage; system or device is down; Customer cannot perform business critical functions and is losing Revenue.

"Severity 2" means a partial site outage/loss of redundancy; a system or component is down; Hotels may be experiencing degradation of service, or loss of resilience.

"Severity 3" means an Incident which is non-business impacting; a hosted system is experiencing minor issues or an individual system component has failed, however is not causing degradation of Service or losing Revenue.

It is assumed that appropriate pro-active monitors are in place at the supplier's data center(s) (Including any on-premises footprint) to alert the supplier of any system malfunction. However, for hotel reported Severity Level 1 incidents, the customer must report the incident to the supplier helpdesk by phone in order to trigger the five-minute response time.

Table 6 General Service Levels

Critical Service Level	Service Level Target	Measurement Window	Measurement Calculation
Percentage of Severity 1 Incidents Responded to On-Time	95% in 5 minutes or less	Monthly	Percentage of Severity 1 Incidents Responded to On-Time = (a) the number of Severity 1 Incidents responded to within 5 minutes or less of the occurrence of the Severity 1 Incident during a Measurement Window, divided by (b) the total number of Severity 1 Incidents that occur during a Measurement Window, expressed as a percentage.
Percentage of Severity 2 Incidents Responded to On-Time	90% in 20 minutes or less	Monthly	Percentage of Severity 2 Incidents Responded to On-Time = (a) the number of Severity 2 Incidents responded to within 20 minutes or less of the occurrence of the Severity 2 Incident during a Measurement Window, divided by (b) the total number of Severity 2 Incidents that occur during a Measurement Window, expressed as a percentage.

Percentage of Severity 3 Incidents Responded to On-Time	90% in 60 minutes or less	Monthly	Percentage of Severity 3 Incidents Responded to On-Time = (a) the number of Severity 3 Incidents responded to within 60 minutes or less of the occurrence of the Severity 3 Incident during a Measurement Window, divided by (b) the total number of Severity 3 Incidents that occur during a Measurement Window, expressed as a percentage.
Percentage of Severity 1 Incidents Resolved On-Time	95% in 4 hour or less	Monthly	Percentage of Severity Level 1 Incidents Resolved On-Time = (a) the number of Severity 1 Incidents Resolved within 4 hours or less of the occurrence of the Severity 1 Incident during a Measurement Window, divided by (b) the total number of Severity 1 Incidents that occur during a Measurement Window, expressed as a percentage.
Percentage of Severity 2 Incidents Resolved On-Time	90% in 12 hours or less	Monthly	Percentage of Severity Level 2 Incidents Resolved On-Time = (a) the number of Severity 2 Incidents Resolved within 12 hours or less of the occurrence of the Severity 2 Incident during a Measurement Window, divided by (b) the total number of Severity 2 Incidents that occur during a Measurement Window, expressed as a percentage.
Percentage of Severity 3 Incidents Resolved On-Time	90% in 24 hours or less	Monthly	Percentage of Severity Level 3 Incidents Resolved On-Time = (a) the number of Severity 3 Incidents Resolved within 24 hours or less of the occurrence of the Severity 3 Incident during a Measurement Window, divided by (b) the total number of Severity 3 Incidents that occur during a Measurement Window, expressed as a percentage.
Hotel Satisfaction	Overall Average of 8 or better	Per Quarter	Hotel Satisfaction Rating = for each customer satisfaction survey, the overall average score must be an 8 or higher.

Table 7 Private Cloud Hosted Services

Critical Service Level	Service Level Target	Measurement Window	Measurement Calculation
Private Cloud Availability (Combined)	100%	Monthly	Private Cloud Availability = (a) total hours in the Measurement Window minus total hours the Private Cloud Service is not Available at both the Primary and Secondary data centers simultaneously in the Measurement Window, divided by (b) total hours in the Measurement Window, expressed as a percentage. For purposes of the Service Levels for Private Cloud Availability, the term "Private Cloud Service" shall require that the following are Available: (i) virtual machine monitors (VMM); (ii) servers; (iii) storage and (iv) Network Connections.
Availability – Private Cloud (Primary Data Center)	99.99% Availability	Monthly	Private Cloud Availability (Primary) = (a) total hours of possible Availability for the Private Cloud Service at the Primary data center in a Measurement Window minus total hours Private Cloud Service is not Available at the primary data center in a Measurement Window, divided by (b) total hours of possible Availability for Private Cloud Service in a Measurement Window, expressed as a percentage. For purposes of the Service Levels for Private Cloud Availability, the term "Private Cloud Service" shall require that the following are Available: (i) virtual machine monitors (VMM); (ii) servers; (iii) storage and (iv) Supplier Connections.
Availability – Private Cloud (Secondary Data Center)	99.99% Availability	Monthly	Private Cloud Availability (Secondary) = (a) total hours of possible Availability for the Private Cloud Service at the Secondary data center in a Measurement Window minus total hours Private Cloud Service is not Available at the Secondary data center in a Measurement Window, divided by (b) total hours of possible Availability for Private Cloud Service in a

			Measurement Window, expressed as a percentage. For purposes of the Service Levels for Private Cloud Availability, the term "Private Cloud Service" shall require that the following are Available: (i) virtual machine monitors (VMM); (ii) servers; (iii) storage and (iv) Supplier Connections.
Capacity Provisioning – Private Cloud	Activate On Demand Capacity for up to a 30% Increase as new hotels are added to system	Per Request	Capacity Provision (Private Cloud) = for each request for additional capacity (new hotel instances) in the Private Cloud Services, the additional capacity requested must be active and deployed when requested.

Table 8 On-Premises Services

Critical Service Level	Service Level Target	Measurement Window	Measurement Calculation
Availability	> 100% (Per Hotel)	Monthly	On-Premises System Availability = (a) for each individual Site, the total hours of possible Availability for the On-Premises Service in a Measurement Window minus the total hours On-Premises Service is not Available in a Measurement Window, divided by (b) total hours of possible Availability for On-Premises Service in a Measurement Window, expressed as a percentage.

Table 9 Managed Hosting Services

Critical Service Level	Service Level Target	Measurement Window	Measurement Calculation
Individual Hotel Instance Availability	99.9%	Monthly	Individual Hotel Instance Availability (a) for each hotel instance, the total hours of possible Availability during a Measurement Window minus total hours that the hotel instance is not Available during a Measurement Window, divided by (b) total hours of possible Availability during a Measurement Window, expressed as a percentage.
Data Backup Service	99.99%	Monthly	Data Protect Backup Service (a) the total hours of possible Availability for the Data Backup Services during a Measurement Window minus total hours that the Data Protect Backup Services are not Available during a Measurement Window, divided by (b) total hours of possible Availability for Data Backup Services during a Measurement Window, expressed as a percentage.