Splunk SOAR Validated Architectures Version 2.0

Product Best Practices

Version 2 on June 3, 2021

Table of Contents

Introduction

Splunk SOAR Validated Architectures (SSVAs) are proven reference architectures for stable, efficient, and repeatable Splunk SOAR deployments. Many of Splunk's existing customers have experienced rapid adoption and expansion, leading to certain challenges as they attempt to scale. New Splunk SOAR customers are increasingly looking for guidelines and certified architectures to ensure that their initial deployment is built on a solid foundation. SSVAs have been developed to help our customers with these growing needs.

Whether you are a new or existing Splunk SOAR customer, SSVAs will help you build an environment that is easier to maintain and simpler to troubleshoot. SSVAs are designed to provide you with the best possible results while minimizing your total cost of ownership. Additionally, your entire Splunk SOAR foundation will be based on a repeatable architecture which will allow you to scale your deployment as your needs evolve over time.

SSVAs offer topology options that consider a wide array of organizational requirements, so you can easily understand and find a topology that is right for your requirements. The Splunk SOAR Validated Architectures selection process will help you match your specific requirements to the topology that best meets your organization's needs. If you are new to Splunk SOAR, we recommend implementing a Validated Architecture for your initial deployment. If you are an existing customer, we recommend that you explore the option of aligning with a Validated Architecture topology. Unless you have unique requirements that make it necessary to build a custom architecture, it is very likely that a Validated Architecture will fulfill your requirements while remaining cost effective.

This white paper will provide you with an overview of SSVAs. Within this whitepaper, you will find the resources you need to go through the SSVA selection process, including the requirements questionnaire, deployment topology diagrams, design principles, and general guidelines.

If you need assistance implementing a Splunk SOAR Validated Architecture, contact Splunk Professional Services (https://www.splunk.com/en_us/support-and-services/splunk-services.html).

Document Structure

SSVAs are broken into three major content areas:

1. Automation and Case Management Tier

Automation and Case Management covers the architecture tier that provides the Splunk SOAR functionality – front-end interface, playbook automation, and data storage.

2. Integrations Tier

The Integrations Tier section guides you in choosing the right integrations to meet your requirements.

3. Design Principles and Best Practices

Design Principles and Best Practices apply to your architecture as a whole and will help you make the correct choices when working out the details of your deployment.

Reasons to Use Splunk SOAR Validated Architectures

Implementing a Validated Architecture will empower you to design and deploy Splunk SOAR more confidently. SSVAs will help you solve some of the most common challenges that organizations face, including:

Performance

- Organizations want to see improvements in performance and stability.

Complexity

- Organizations sometimes run into the pitfalls of custom-built deployments, especially when they have grown too rapidly or organically. In such cases, unnecessary complexity may have been introduced into the environment. This complexity can become a serious barrier when attempting to scale.

Efficiency

- To derive the maximum benefits from the Splunk deployment, organizations must improve the efficiency of operations and accelerate time to value.

Cost

- Organizations are seeking ways to reduce total cost of ownership (TCO), while fulfilling all of their requirements.

Agility

- Organizations will need to adapt to change as they scale and grow.

Maintenance

- Optimization of the environment is often necessary in order to reduce maintenance efforts.

Scalability

- Organizations must have the ability to scale efficiently and seamlessly.

Verification

- Stakeholders within the organization want the assurance that their Splunk deployment is built on best practices.

1.1 Pillars of Splunk SOAR Validated Architectures

Splunk SOAR Validated Architectures are built on the following foundational pillars. For more information on these design pillars, refer to Appendix "A" below.

AVAILABILITY	***PERFORMANCE***	***SCALABILITY***	***SECURITY***	***MANAGEABILITY***
*The system meets customer **continuity requirements** is able to recover from planned and unplanned outages or disruptions.*	*The system can **maintain an optimal level of service** under varying usage patterns.*	*The system is designed to scale on all tiers, allowing you to **handle increased workloads effectively**.*	*The system is **designed to protect data, configurations, and assets** while continuing to deliver value.*	*The system is **centrally operable and manageable across all tiers**.*

These pillars are in direct support of the **Platform Management & Support** Service in the Splunk Center of Excellence model.

1.2 What to Expect from Splunk SOAR Validated Architectures

Please note that SSVAs do not include deployment technologies or deployment sizing. The reasoning for this is as follows:

- Deployment technologies, such as operating systems and server hardware, are considered implementation choices in the context of SSVAs. Different customers will have different choices, so a generalization is not easily possible.
- Deployment sizing requires an evaluation of data ingest volume, data types, file volumes, and playbook use cases, which tend to be very customer-specific and generally have no bearing on the fundamental deployment architecture itself.
- Customer should engage with Professional Services to ensure that proper sizing and adequate loading is configured. Currently, there is no formal sizing guides for Splunk SOAR. There is a recommended single instance sizing for productions systems located here for on premise systems: https://docs.splunk.com/Documentation/Phantom/latest/Install/ProdRequirements

SSVAs <u>will</u> provide:	***SSVAs do <u>not</u> provide:***
✓ *Clustered and non-clustered deployment options.* ✓ *Diagrams of the reference architecture.* ✓ *Guidelines to help you select the architecture that is right for you* ✓ *Tier-specific recommendations.* ✓ *Best practices for building out your Splunk SOAR deployment*	✗ *Implementation choices (OS, bare metal vs. virtual vs. Cloud etc.).* ✗ *Deployment sizing.* ✗ *A prescriptive approval of your architecture. Note: SSVAs provide recommendations and guidelines, so you can ultimately make the right decision for your organization.* ✗ *A topology suggestion for every possible deployment scenario. In some cases, unique factors may require a custom architecture to be developed. Splunk experts are available to help with any custom solutions you need. If you are an existing customer, reach out to your Splunk Account Team. If you are new to Splunk, you can reach us here (https://www.splunk.com/en_us/talk-to-sales.html).*

1.3 Roles and Responsibilities

Splunk SOAR Validated Architectures are highly relevant to the concerns of decision makers and administrators. Architects, consultants, Splunk SOAR administrators, and managed service providers should all be involved in the SSVA selection process. You will find a description of each of these roles below:

Role	**Description**
Architects	*Responsible for architecting Splunk deployments to meet enterprise needs.*
Consultants	*Responsible for providing services for Splunk architecture, design, and implementation.*
Splunk SOAR Specialists	*Responsible for managing the Splunk SOAR product lifecycle.*
Managed Service Providers	*Entities that deploy and run Splunk as a service for customers.*

1.4 Overview of the Splunk SOAR Validated Architectures Selection Process

The Splunk SOAR Validated Architectures selection process will help you identify the simplest and most streamlined architecture that meets all of your organization's needs.

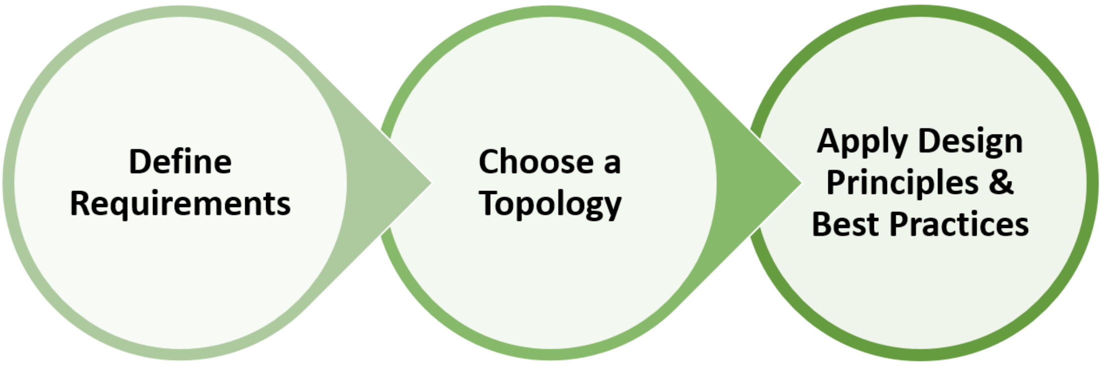

Steps in the Selection Process	***Goals***	***Considerations***
Step 1: Define Requirements for: ***a) Automation and Case Management*** ***b) Integration Needs***	*Define requirements.*	● *Decision-makers, stakeholders, and admins should collaborate to identify and define your organization's requirements.* ● *If you already have a deployment in place, you can evaluate your current architecture to see what it would take to move to a validated model.* *For a questionnaire that will help you define your requirements, refer to Step 1.5 below.*
Step 2: Choose a Topology for: ***Automation and Case Management***	*Choose a topology that meets identified requirements.*	● *You'll choose a topology that best meets your requirements.* ● *Keep things simple and in accordance with the SSVA, so you can appreciate the easier path to scalability.* *For diagrams and descriptions of topology options, refer to Step 2 below.*
Step 3: Apply Design Principles and Best Practices	*Prioritize your design principles and review tier-specific implementation best practices.*	● *Each design principle reinforces one or more of the pillars of Splunk SOAR Validated architectures.* ● *You'll prioritize design principles in accordance with the needs of your organization.* ● *Tier-specific recommendations will guide your topology implementation.* *For a breakdown of design principles, refer to Step 3 below.*

1.5 Step 1: Define Your Requirements for Automation

To select the appropriate deployment topology, you will need to do a deep dive into your requirements. Once you have defined your requirements you will be able to choose the simplest, most cost-effective way to deploy Splunk SOAR. Below you will find a questionnaire that will help you define key requirements areas for the indexing and search tiers of your deployment.

The requirements questionnaire focuses on areas that will have a direct impact on your deployment topology. Therefore, we highly recommend that you record your answers to the questions below before choosing a topology in the next step.

1.5.1 Things to Keep Under Consideration

Review your use cases

*Headles*s operation is considered only for design and playbook execution and very minimal if any customer interaction other than design and administrative operations. *Case Management* operation is considered full-use functionality of the platform and is measured by the number of concurrent customers using the platform. As you define your requirements, you should think about the intended usage of Splunk SOAR. For example, the topology for a "headless" deployment of Splunk SOAR acting as an automation backend will require far fewer interactive users than a deployment in which Splunk SOAR will be the case management system of record. You should fully consider use cases involving:

- Headless vs Case management operations
- Reporting
- Availability
- Disaster Recovery Requirements
- Other use case scenarios specific to your organization

Depending on your use case scenarios, your deployment may need to provide additional architectural characteristics.

Think about future growth

You will need to think about your immediate needs in order to define your requirements. However, you should also consider future growth and scalability. Scaling your deployment may require expenditures, additional staffing, or other resources you may want to start planning for today.   We have started this plan for the average use for a customer to able to last at least 1 year of operations before modification should be considered. This will dependent on the usage, volume, playbook and asset configurations. Since these vary greatly between customers, please understand that this planning varies between customers and is an estimate.

1.5.2 Topology Categories

The following is a key to SSVA topology categories. These categories are used in the questionnaire below. You will also find references to these categories in the next steps of the SSVA selection process.

Platform Topology Categories

Category Code	Explanation
S	*Category "S" indicates a single-server Splunk SOAR deployment*
X	*Category “X” indicates an externalized single-instance Splunk SOAR deployment with externalized shared services (e.g., file share, database)*
D	*Category "D" indicates a distributed Splunk SOAR deployment across two sites configured in Warm Standby Mode. Warm Standby is not supported with externalized shared services.*
C	*Category "C" indicates the need for a clustered automation and case management tier (data replication and high capacity for automation is required). Clustered SOAR deployments always require externalized shared services.*

Instance Requirements Categories

Category Number	Explanation
0	*Category "0" indicates a Software as a Service model and no physical hardware is required.*
1	*Category "1" indicates a single or up to 3 hardware/virtual instances* will meet requirements*
2	*Category "2" indicates the need for multiple hardware/virtual instances minimum of 3 and up to 8 separate hardware instances per site*

** Instance is defined as a single-instance appliance (virtual or on physical hardware), single-instance with external shared components, a SaaS tenant, or a single SOAR cluster*

Integration Tier Categories

Category Code	Explanation
E	*Category “E” indicates that SOAR will leverage an external Splunk instance for custom reporting capability*
E+	*Category “E+” indicates that SOAR will leverage one or all of the optional customer provided load balancers, NFS, or utilize AWS Services*
CE	*Category “CE” indicates that SOAR will leverage an external Splunk Cloud Enterprise instance for custom reporting capability or ingestion source.*
CE+	*Category “CE+” indicates that SOAR will leverage an external Splunk Cloud Enterprise instance for custom reporting capability and AWS Services*

1.5.2.1 Define Your Requirements for Automation

♦ See the key above for an explanation of topology category codes. If you answer "yes" to multiple questions, use the topology category code for the highest numbered question.

#	Question	Considerations	Impact on Topology	Core Platform Topology	Instance Requirements
1	*Are you wanting a solution that reduces your infrastructure & maintenance costs?*	*Consideration for case management or cloud-based automation.* *High volume > 750 events per hour and High-performance transactions is* ***not*** *required.* *Supports up to ~20 concurrent users*	*Candidate for SaaS solution that provides reduced maintenance costs and managed infrastructure.* *Certain on-premises capabilities are not present.* *Constrained to regional deployments*	*S*	*0*
2	*Is your expected daily data ingest up to* ***500 forwarded events/hour****?* *These are forwarded events to SOAR and* ***not*** *ES or Splunk EPS calculations*	*Considered short-term growth in the daily ingest (~6-12 month)* *These are forwarded events to SOAR and* ***not*** *ES or Splunk EPS calculation*	*Candidate for a single server deployment, depending on answers to availability-related questions*	*S*	*1*
3	*Are you ingesting up to* ***500 events/hour*** *and using case management with* ***less than 10 personnel continuously****?*	*Need to consider long-term growth of database and file collections (~6-12 month)* *If you have your own site replication capabilities like VMware Site Recovery Manager or in AWS region and don't require multi-regional support*	*Candidate for externalizing shared services like file and database to increase performance and local redundancy.*	*X*	*1*
4	*Do you require Disaster Recovery for data ingestion or automation?*	*If you are planning on using SOAR for data enrichment only and/or batched automation jobs, an interruption of service may be acceptable.* *This is the most common customer use configuration. It provides* ***Active/Passive Warm Standby services****.*	*Requires two SOAR instances deployed in a warm standby configuration to support continuous ingest and automation*	*D*	*1*

#	Question	Considerations	Impact on Topology	Core Platform Topology	Instance Requirements
5	*Is your expected daily data ingest greater than* ***~500 forwarded events/hour****?* *These are forwarded events to SOAR and* ***not*** *ES or Splunk EPS calculations*	*Automated playbook development greatly affects event ingestion. Default build can handle ingestion of ~10,000 events a day.* *This configuration is considered high volume and high availability. It is a horizontally scalable deployment.* *If your continuous user count is above 10 users a day, you should consider a clustered deployment*	*Requires a High-Capacity Clustered Automation and Case Management tier with externalized shared services.*	*C*	*1*
6	*Assuming an available Splunk SOAR instance: Does your data need to* ***ensure active/active availability for automation execution?***	*If your use case is swiftly, responding to, and remediating issues from ingested alerts, automated blocking or removal and local site down time is not acceptable.*	*Requires a High-Capacity Clustered Automation and Case Management tier or a Distributed system with Customer provided equipment.* *D is TCO is less than C TCO. With D customer needs to provide adequate application-level load balancers and must script the process to fail over.*	*D/C*	*1*
7	*Do you operate multiple data centers and require recovery of your Splunk SOAR environment in case of a data center outage?*	*Disaster recovery requirements may dictate continuous prescribe RTO/RPO goals for manual disaster recovery*	*Failover will require two SOAR platforms operating in in warm standby mode.* *Use "D" if SOAR* ***is*** *being used for case management* *Use "M" if SOAR i****s not*** *being used for case management*	*D/M*	*2*
8	*Do you need to support many concurrent users?*	*Requirements for more than ~10 concurrent users typically require horizontal scaling of the Automation and Case Management Tier*	*Requires a clustered Automation and Case Management tier with externalized shared services.*	*C*	*1*

1.5.2.2 Questionnaire 1: Define your requirements for integrations:

#	Question	Considerations	Impact on Topology	Integration Requirements
1	*Do you require SOAR clustering (see above questionnaire)?*		*Clustering requires the externalization of Splunk to act as the SOAR Reporting Tier*	*E*
2	*Do you require the ability to create custom dashboards and reports for SOAR metrics?*	*Standard reports & metrics can be produced without any external reporting components or resources.* *Indicate a yes here, if you have custom metrics you want to measure or detailed reporting requirements.*	*Custom reporting and dashboards require an external Splunk Enterprise instance to act as the SOAR Reporting Tier*	*E*
3	*Do you want to provide your own* ***external infrastructure*** *or utilize cloud infrastructures*	*If the customer wants to provide separate file services or database services.*	*Reduced cost of ownership*	*E+*
4	*Do you require the ability to integrate* ***Splunk Cloud Infrastructure*** *with customer provided infrastructure (on premise)*	*If you are using Splunk Cloud or Cloud ES and yet want SOAR on-premises only.*	*Requires the use of Splunk SOAR in the DMZ or VPC connections to provide cloud to premise connectivity.*	*CE*
5	*Do you require the ability to integrate* ***Splunk Cloud Infrastructure*** *with customer provided or other Cloud infrastructure*	*If the customer wants to provide separate file services or database services.* *If you are using Splunk Cloud or Cloud ES and yet want to integrate SOAR*	*Custom reporting and dashboards require an external Splunk instance to act as the SOAR Reporting Tier*	*CE+*

1.5.3 How to Determine Your Topology Category Code

Based on your answers to the requirements questionnaire above, you will end up with a combined topology category indicator that will allow you to identify the best topology for your needs. Instructions and examples are provided below.

Instructions

1. Write down the questions to which you answered "yes".
2. If you answered "yes" to multiple questions, follow the topology recommendation for the highest numbered question. If you see multiple topology options (for example, "S/C"), look at the previous questions to determine which option is best suited for you.
3. Your reporting tier code will be the letter representing the question(s) to which you answered yes. If your answer results are yes, then use E for your reporting tier code.

Example #1

Let's say you answered "yes" to questions #4, #5 and #7, and you need custom reporting requirements (1b. Question #2). You will end up with a topology category of "C1E", indicating the need for a clustered indexing tier with an external reporting server.

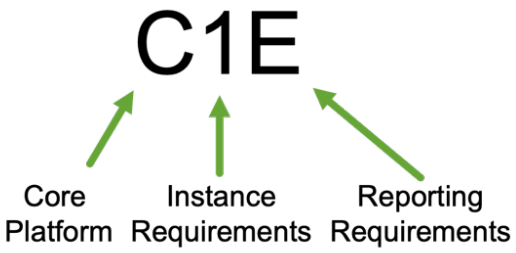

Example #1

Let's say you answered "yes" to questions #4, #5 and #7, and you need custom reporting requirements (1b. Question #2). You will end up with a topology category of "C1E", indicating the need for a clustered indexing tier with an external reporting server(s).

1.6 Step 2: Choose a Topology for Automation

Topologies are generally split into non-clustered and clustered deployments. Non-clustered deployments require the least number of distinct components and have excellent scalability characteristics.

Remember: The primary goal of the SSVA selection process is to allow you to build what you need without introducing unnecessary components.

Note

While you may choose to implement a topology that provides additional benefits beyond your immediate needs, keep in mind that this will likely result in unnecessary costs. Moreover, the introduction of additional complexity is often counter-productive to operational efficiency.

Important Note about topology diagrams

*The icons in the topology diagrams represent **functional Splunk roles** and do not imply dedicated infrastructure to run them. Please see the Appendix for guidance as to which Splunk roles can be collocated on the same infrastructure/server.*

1.6.1 Using Your Topology Category Code

Before selecting a topology option, it is highly recommended that you complete the requirements questionnaire to determine your topology category code. If you have not yet done this, please go back and complete the previous step above. Once you have your topology category code you will be able to identify the deployment option that is the best fit for your stated requirements. These Automation tier topologies have been matched and aligned to your recommended Splunk Index and Search topologies.

1.6.2 Non-clustered deployment options

Below you will find the following the available topology options:

Type of Deployment	Topology Category Code(s)
Single Tenant Deployment using Software as a Service with Splunk managed infrastructure	*S0*
Single Tenant Deployment using Software as a Service with Splunk managed infrastructure with on premise integration	*S0E*
Single Tenant Deployment using Software as a Service with Splunk managed infrastructure with hybrid cloud solution	*S0E+*
Single Tenant Deployment using Software as a Service with Splunk managed infrastructure with Splunk Cloud Integration	*S0CE*
Single Server Deployment using embedded Splunk integration - Single site	*S1*
Single Server Deployment using Splunk Enterprise Integration - Single site	*S1E*
Single Server Deployment using external Shared Services with embedded Splunk	*X1*
Single Server Deployment using external Shared Services with Splunk Enterprise Integration	*X1E*
Distributed Warm Standby Deployment using embedded Splunk integration - Multiple site	*D2*
Distributed Warm Standby Deployment using embedded Splunk with customer provided infrastructure – Multiple site	*D2E*
Distributed Warm Standby Deployment using embedded Splunk with customer provided infrastructure – Multiple site	*D2E+*
Distributed Warm Standby Deployment with Splunk Cloud Integration - Multiple site	*D2CE*
Distributed Warm Standby Deployment with Splunk Cloud Integration - Multiple site	*D2CE+*
High Capacity Clustered Deployment - Single Site	*C1*
High Capacity Clustered Deployment - Single Site	*C1E*
High Capacity Clustered Deployment with AWS Integrations or customer provided infrastructure – Single Site	*C1E+*
High Capacity Clustered Deployment with Splunk Cloud – Single Site	*C1CE*
High Capacity Clustered Deployment with Splunk Cloud and AWS Integrations – Single Site	*C1CE+*
High Capacity Clustered Deployment - Multiple Site	*M2E*
High Capacity Clustered Deployment with Splunk Cloud - Multiple Site	*M2CE*
High Capacity Clustered Deployment with Splunk Cloud and AWS Integrations - Multiple Site	*M2CE+*

1.6.2.1 Single Server Deployment (S0)

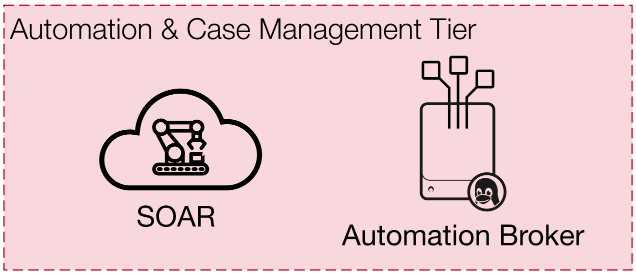

For an explanation of topology components, refer to Appendix "B" below.

Description of the Single Server Deployment (S0)	Limitations
This deployment topology provides you with a very cost-effective solution if your environment meets all the following criteria: *a) you have requirements to provide high-availability or automatic disaster recovery for your Splunk SOAR Deployment,* *b) your daily event ingestion is < 750 events per day, and* *c) you have approximately 20 concurrent users in your environment* *d)* ***suitable for a production environment*** *The primary benefits of this topology include easy manageability, good search performance for smaller ingest and concurrent user count.* *This topology is commonly used in production environments and the primary benefits of this topology include easy manageability and a fixed TCO.*	● *Scalability provided by Splunk managed infrastructure* ● *Constrained to regional deployments* ● *Limited to Internet accessible integrations or Automation Broker integrations*

1.6.2.2 Single Server Deployment (S0E)

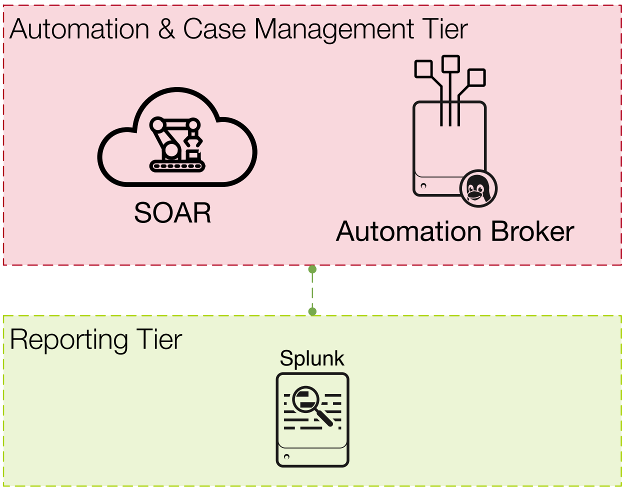

For an explanation of topology components, refer to Appendix "B" below.

Description of the Single Server Deployment (S0E)	Limitations
This deployment topology provides you with a very cost-effective solution if your environment meets all the following criteria: *a) you have requirements to provide high-availability or automatic disaster recovery for your Splunk SOAR Deployment,* *b) your daily event ingestion is < 750 events per day, and* *c) you have approximately 20 concurrent users in your environment* *d)* ***suitable for a production environment*** *The primary benefits of this topology include easy manageability, good search performance for smaller ingest and concurrent user count.* *Common integration with SaaS and hybrid cloud integrations* *This topology is commonly used in production environments and the primary benefits of this topology include easy manageability and a fixed TCO.*	● *Scalability provided by Splunk managed infrastructure* ● *Ingestion and reporting are mostly provided by Splunk Integrations* ● *Splunk forwarding will need access to the public internet* ● *Constrained to regional deployments*

1.6.2.3 Single Server Deployment (S0E+)

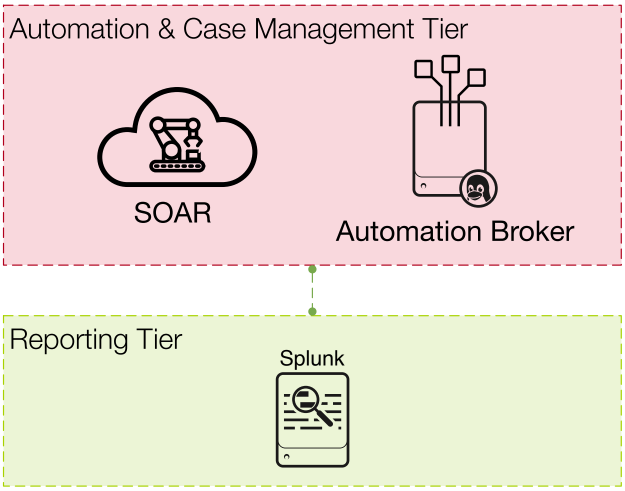

For an explanation of topology components, refer to Appendix "B" below.

Description of the Single Server Deployment (S0E+)	Limitations
This deployment topology provides you with a very cost-effective solution if your environment meets all the following criteria: *a) you have requirements to provide high-availability or automatic disaster recovery for your Splunk SOAR Deployment,* *b) your daily event ingestion is < 750 events per day, and* *c) you have approximately 20 concurrent users in your environment* *d)* ***suitable for a production environment*** *The primary benefits of this topology include easy manageability, good search performance for smaller ingest and concurrent user count.* *Common integration with SaaS and hybrid cloud integrations* *This topology is commonly used in production environments and the primary benefits of this topology include easy manageability and a fixed TCO.*	● *Scalability provided by Splunk managed infrastructure* ● *Ingestion and reporting are mostly provided by Splunk Integrations* ● *Splunk forwarding will need access to the public internet* ● *Constrained to regional deployments*

1.6.2.4 Single Server Deployment (S0CE)

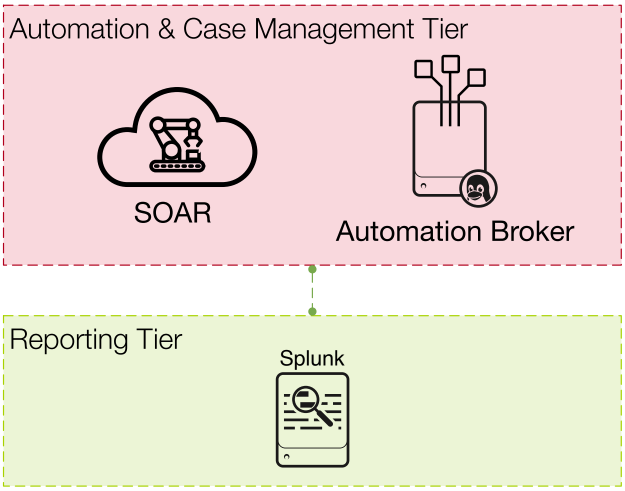

For an explanation of topology components, refer to Appendix "B" below.

Description of the Single Server Deployment (S0CE)	Limitations
This deployment topology provides you with a very cost-effective solution if your environment meets all the following criteria: *a) you have requirements to provide high-availability or automatic disaster recovery for your Splunk SOAR Deployment,* *b) your daily event ingestion is < 750 events per day, and* *c) you have approximately 20 concurrent users in your environment* *d) **suitable for a production environment*** *The primary benefits of this topology include easy manageability, good search performance for smaller ingest and concurrent user count.* *Common integration for SaaS to SaaS Splunk Components* *This topology is commonly used in production environments and the primary benefits of this topology include easy manageability and a fixed TCO.*	● *Scalability provided by Splunk managed infrastructure* ● *Ingestion and reporting are mostly provided by Splunk Cloud Integrations* ● *Constrained to regional deployments*

1.6.2.5 Single Server Deployment (S1)

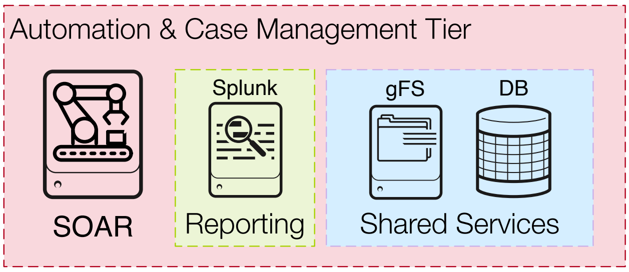

For an explanation of topology components, refer to Appendix "B" below.

Description of the Single Server Deployment (S1)	Limitations
This deployment topology provides you with a very cost-effective solution if your environment meets all the following criteria: *a) you do not have any requirements to provide high-availability or automatic disaster recovery for your Splunk SOAR Deployment,* *b) your daily event ingestion is < ~100 events per day, and* *c) you have a small number of users.* *d) **suitable for development environment*** *The primary benefits of this topology include easy manageability, good search performance for smaller ingest and concurrent user count.* *This topology is commonly used in development environments and the primary benefits of this topology include easy manageability and a fixed TCO.*	• *No High Availability* • *Scalability limited by hardware capacity* • *Reporting limited to standard SOAR reports and/or API usage*

1.6.2.6 SOAR Server Deployment (S1E)

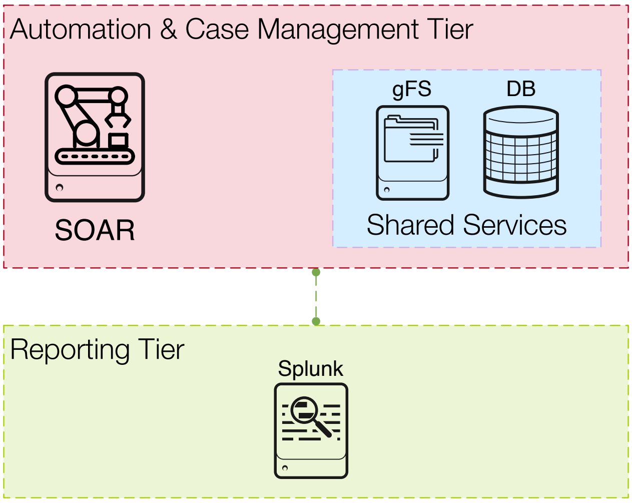

For an explanation of topology components, refer to Appendix "B" below.

Description of the Single Server Deployment (S1E)	**Limitations**
This deployment topology provides you with a very cost-effective solution if your environment meets all of the following criteria: *a) you do not have any requirements to provide high-availability or automatic disaster recovery for your Splunk SOAR Deployment,* *b) your daily event ingestion is < ~100 events per hour, and* *c) you have a small number of users.* *d)* ***suitable for development environment and for developing external SOAR reporting*** *Externalizing the reporting tier facilitates the creation of dashboards and custom reports in a separate Splunk instance.* *This topology should be leverage for small environment that want a simple deployment but still want the robust and flexible reporting capabilities that an external reporting tier can provide.*	● *No High Availability for ingestion/automation/reporting* ● *Scalability limited by hardware capacity*

1.6.2.7 SOAR Single Server deployment with external Shared Services and Internal Reporting (X1)

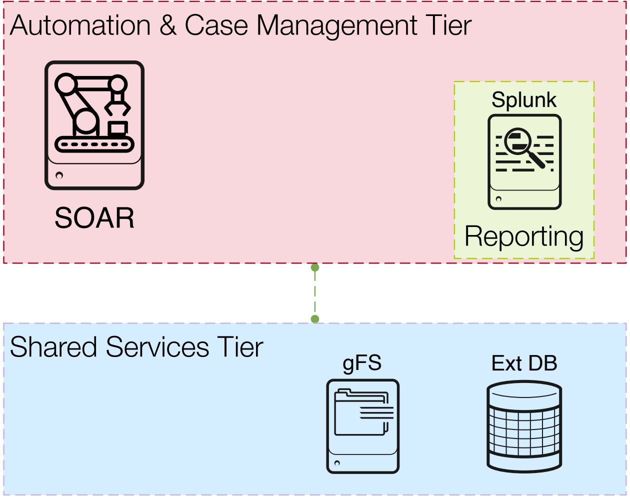

For an explanation of topology components, refer to Appendix "B" below.

Description of the SOAR Single Server deployment with external Shared Services and Internal Reporting (X1)	**Limitations**
This deployment topology provides you with a very cost-effective solution if your environment meets all of the following criteria: *a) you do not have any requirements to provide high-availability or automatic disaster recovery for your Splunk SOAR deployment,* *b) your daily data ingest is over ~100 events/hour, and* *c) you have a less than 10 users with non-critical use cases.* *This topology is typically used for smaller, non-business-critical use-cases (often departmental in nature).* *By externalizing the shared services, some additional flexibility in scaling is provided by distributing load across multiple services.*	● *No High Availability for ingestion/automation* ● *Automation scalability limited by hardware capacity* ● *Reporting limited to standard SOAR reports and/or API usage*

1.6.2.8 SOAR Single Server deployment with external Shared Services and Internal Reporting (X1E)

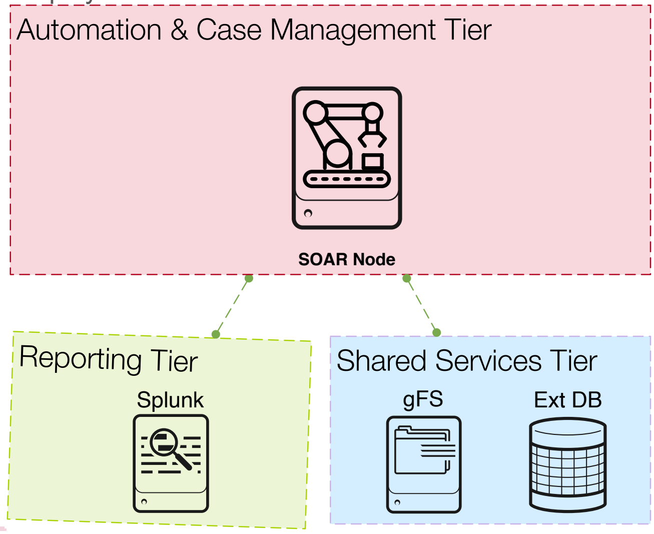

For an explanation of topology components, refer to Appendix "B" below.

Description of the SOAR Single Server deployment with external Shared Services and Internal Reporting (X1E)	**Limitations**
This deployment topology provides you with a very cost-effective solution if your environment meets all of the following criteria: *a) you do not have any requirements to provide high-availability or automatic disaster recovery for your Splunk SOAR deployment,* *b) your daily data ingest is over ~100 events/hour, and* *c) you have a less than 10 users with non-critical use cases.* *This topology is typically used for smaller, non-business-critical use-cases (often departmental in nature).* *By externalizing the shared services, some additional flexibility in scaling is provided by distributing load across multiple services.* *Externalizing the reporting tier facilitates the creation of dashboards and custom reports in a separate Splunk instance.*	• *No High Availability for ingestion/automation* • *Automation scalability limited by hardware capacity*

1.6.2.9 Distributed SOAR Deployment with Warm Standby and Internal Reporting (D1)

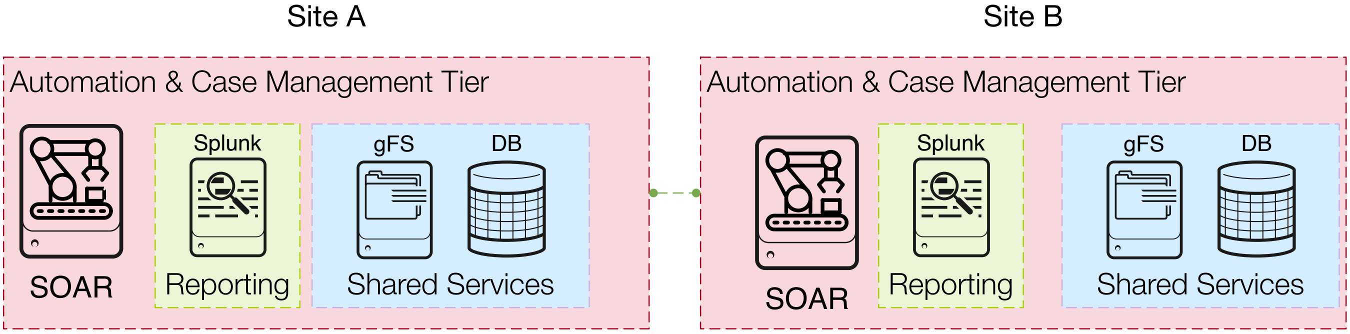

For an explanation of topology components, refer to Appendix "B" below.

Description of Distributed SOAR Deployment with Warm Standby and Internal Reporting (D1)	Limitations
This architecture is suitable for most organizations requiring full disaster recovery with high availability and multi-regional support. *Automation tier may need to move to a more resilient topology for environments where recovering from disaster scenarios is of high importance. In this instance, the SOAR Warm Standby Deployment is recommended. This architecture maintains the simplicity of keeping all of the SOAR services contained on one server while providing an additional instance for failover in the event of a primary outage disaster with recovery in minutes.* *This topology is commonly used in production environments and the primary benefits of this topology include easy manageability and a fixed TCO.*	● *Limited High Availability* ● *Scalability limited by hardware capacity* ● *Reporting limited to standard SOAR reports and/or API usage*

1.6.2.10 Distributed SOAR Deployment with Warm Standby and Internal Reporting (D2E)

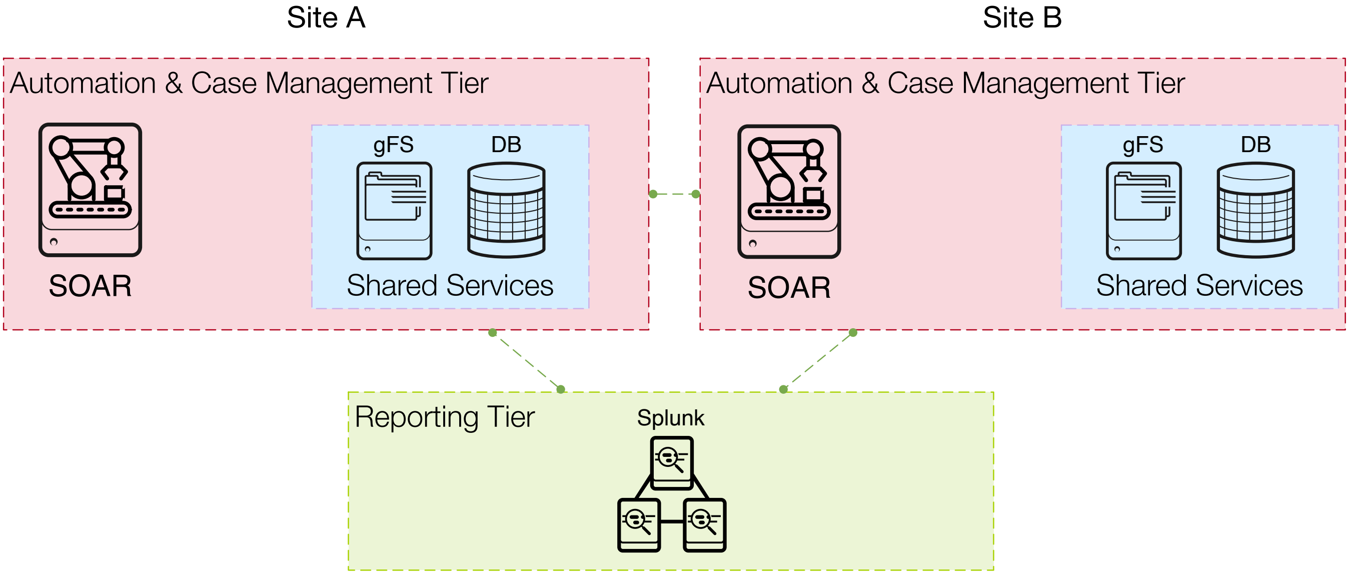

For an explanation of topology components, refer to Appendix "B" below.

Description of Distributed SOAR Deployment with Warm Standby and External Reporting (D2E)	Limitations
*This architecture is **best suitable for most organizations** requiring full disaster recovery with high availability and multi-regional support. This supports customers using Splunk SOAR case management capabilities.* *Automation tier may need to move to a more resilient topology for environments where recovering from disaster scenarios is of high importance. In this instance, the SOAR Warm Standby Deployment is recommended. This architecture maintains the simplicity of keeping all of the SOAR services contained on one server while providing an additional instance for failover in the event of a primary outage or site-level disaster with recovery in minutes.* *Optional configurations with a customer provided load balancer will provide automated failover within seconds.* *This automation tier topology is commonly used in production environments and the primary benefits of this topology include easy manageability and a fixed TCO.* *Externalizing the reporting tier facilitates the creation of dashboards and custom reports in a separate Splunk instance.*	● *Limited High Availability* ● *Scalability limited by hardware capacity*

1.6.3 Clustered deployment options

Below you will find the following topology options:

Type of Deployment	Topology Category Code(s)
Clustered SOAR Deployment	*C1E*

1.6.3.1 High Capacity Clustered Deployment - Single Site (C1E)

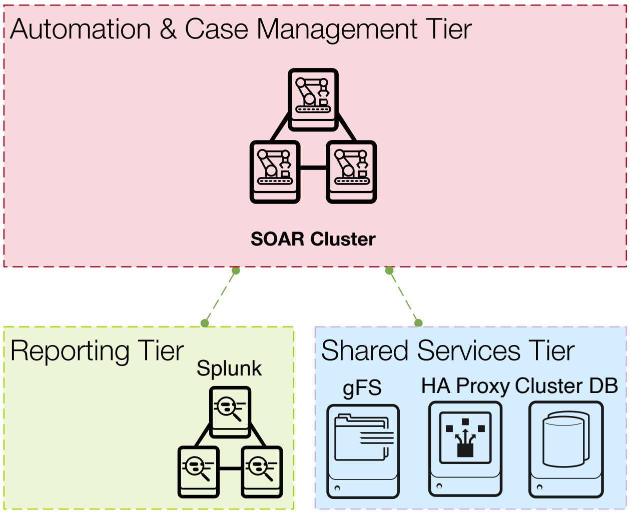

For an explanation of topology components, see Appendix "B" below.

Description of High Capacity Clustered Deployment - Single Site (C1E)	Limitations
This architecture is suitable for organizations that need High Capacity and high availability for actions and ingestion. Customers should have an independent regional replication architecture with a recovery time objective of hours and minutes. *This topology introduces automation high available and high capacity processing in conjunction with **no failover processes**. This provides high availability of data in case of automation peer node, database, and file services failures. However, you should be aware that this applies only to the automation tier and does not protect against search head failure or provide for customized reporting.* *A high availability proxy is provided that ensures ensure proper load balancing of users across the site.* ***Note:** If your category code is C/M1 (i.e., you intend to deploy Splunk SOAR without any external Splunk instance), a **single dedicated** search head is deployed with the appliance and will need to be externalized also. However, this will only impact global searching on the platform.*	- *No High Availability for SOAR platform search* - *No automatic DR capability in case of data center outage*

1.6.3.2 Distributed Clustered Deployment - Multiple Site (M2E)

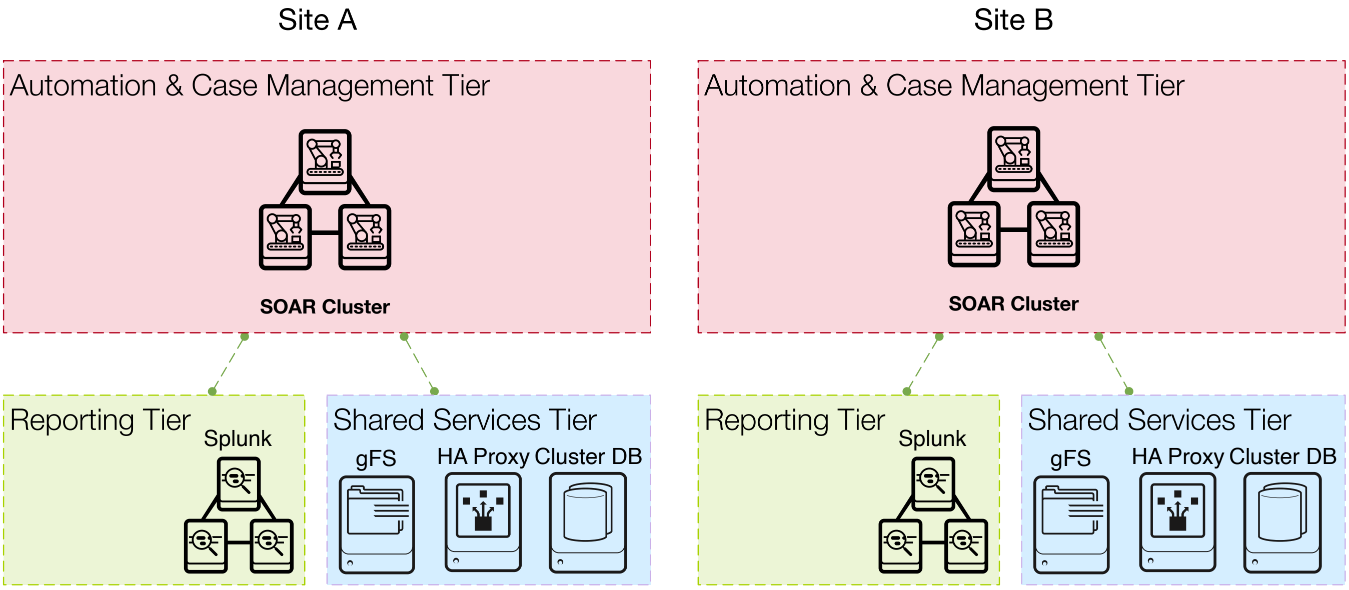

For an explanation of topology components, see Appendix "B" below.

Description of Distributed Clustered Deployment - Single Site (M2E)	Limitations
This architecture is best suitable for organizations that require High Capacity and Availability requirements and don't have data persistence or case management requirements for Splunk SOAR and require near instantaneous recovery time objectives. *This topology introduces automation high available and high capacity processing in conjunction with regional **failover processes**. This provides high availability of data in case of automation peer node, database, and file services failures. However, you should be aware that this applies only to the automation tier and does not protect against search head failure or provide for customized reporting.* ***Customer provided load balancers** and CICD procedures can provide automated failover with necessary application and playbook replication between sites.* ***Note: No event data will be replicated between sites.** Distributed Splunk core architectures can support consistent and multi-site reporting and data visibility.*	● *No automatic replication between sites* ● *No automatic DR capability in case of data center outage* ● *Limited HA with provided Splunk components*

1.6.3.3 Distributed Clustered Deployment - Multiple Site (M2CE)

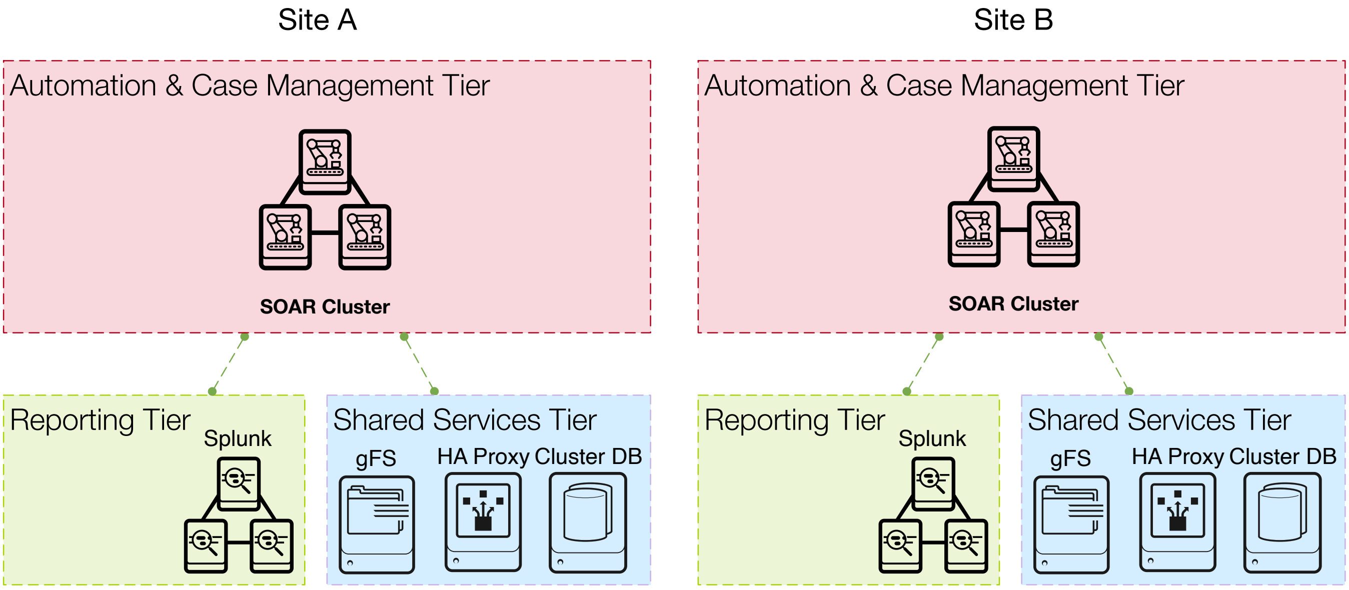

For an explanation of topology components, see Appendix "B" below.

Description of Distributed Clustered Deployment – Multiple Site (M2CE)	Limitations
This architecture will require consultation with your ***Splunk Architect for a custom solution.*** *There are workarounds for limitations with Splunk and customer provide infrastructure*	- *No automatic replication between sites* - *No automatic DR capability in case of data center outage*

1.6.3.4 Distributed Clustered Deployment - Multiple Site (M2CE+)

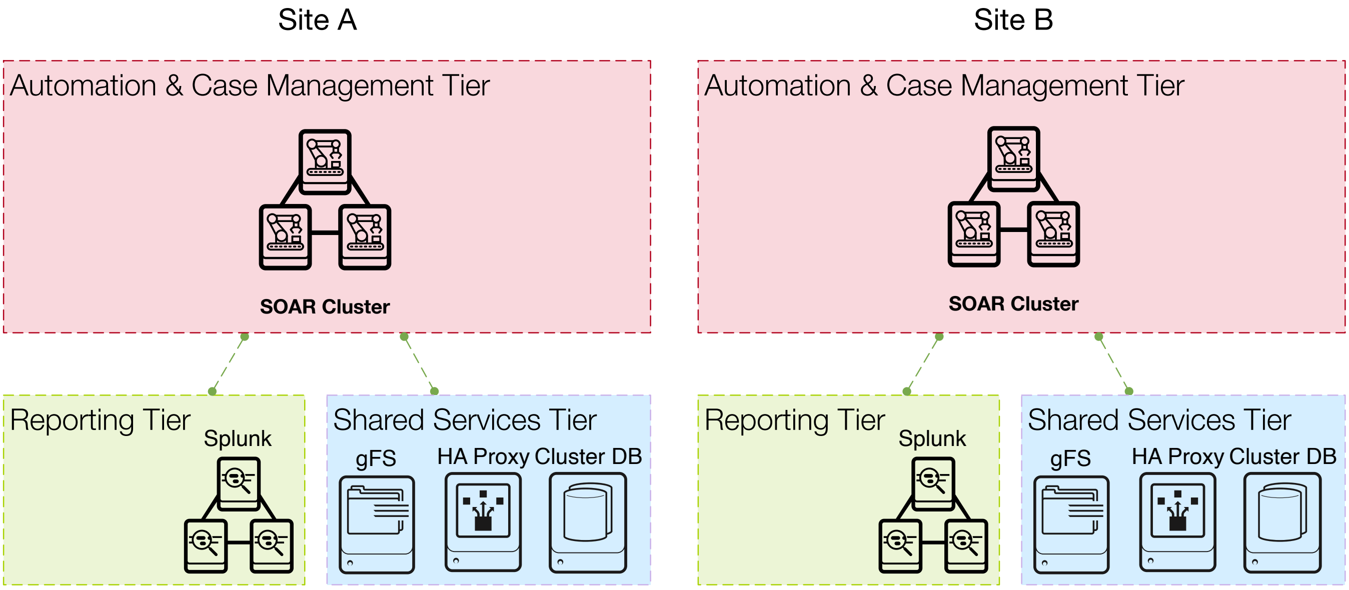

For an explanation of topology components, see Appendix "B" below.

Description of Distributed Clustered Deployment – Multiple Site (M2CE+)	Limitations
*This architecture will require consultation with your **Splunk Architect for a custom solution.*** *There are workarounds for limitations with Splunk and customer provide infrastructure*	- *No automatic replication between sites* - *No automatic DR capability in case of data center outage*

1.7 Step 3: Apply Design Principles and Best Practices

1.7.1 Deployment and Integrations Architectural Diagrams

Below you will find architectural designs and best practices separated by deployment model.

SSVA architectural designs cover all of the following deployment tiers and integrations.

Tier	***Definition***
Automation and Case Management	● *SOAR Node(s)*
Shared Services	● *HA Proxy / Network Load Balancer* ● *File Server* ● *Clustered Database (postgreSQL only)*
Integrations	● *Splunk Enterprise* ● *Splunk Cloud (and Cloud ES)* ● *AWS Cloud Services* ● *Cloud Integrations*

S0 - Single Tenant Deployment using Software as a Service with Splunk managed infrastructure

Intentionally left blank for diagram below

SpVA: S0

Automation Broker Sizing	Min	4 Cores	8 GBs
	Recommended	8 Cores	16 GBs

SOAR Cloud region	Cloud Gateway service region
us-west-2	us-east-1
us-east-1	us-east-1
ca-central-1	us-east-1
eu-west-1	eu-central-1
eu-west-2	eu-central-1
eu-west-3	eu-central-1
ap-southeast-1	ap-southeast-2
ap-southeast-2	ap-southeast-2
ap-northeast-1	ap-southeast-2
ap-northeast-2	ap-southeast-2

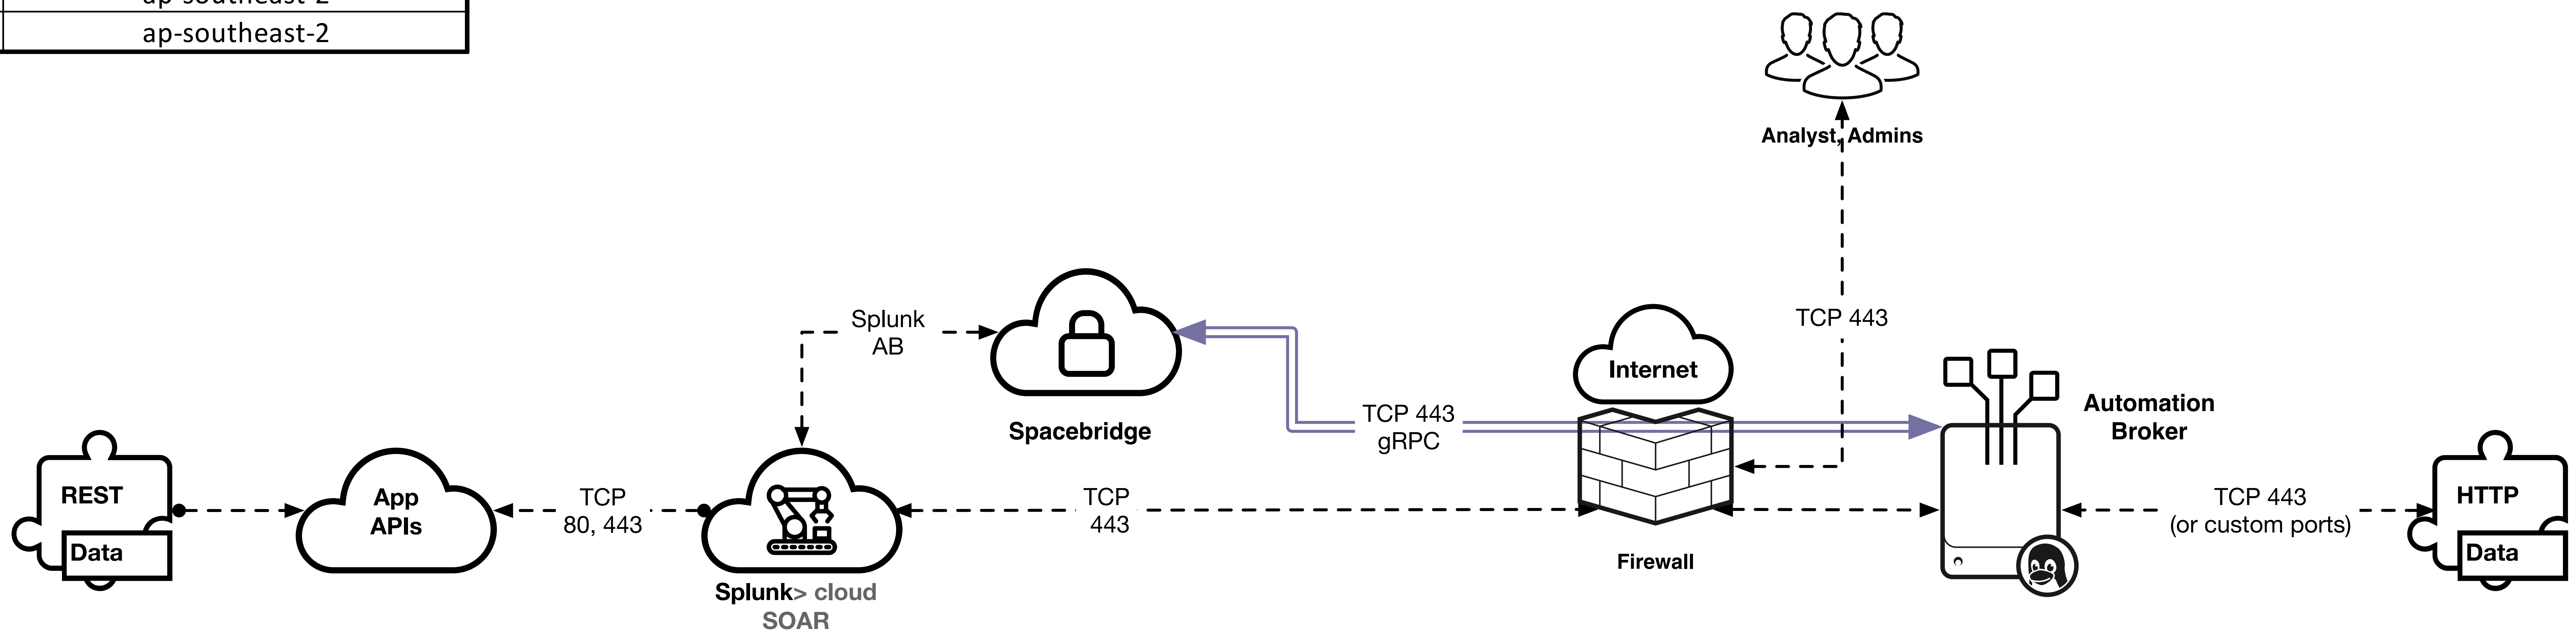

PROS:
Least Administrative overhead
Infrastructure and upgrades managed for customer

CONS:
Performance constrained by number of users

Automation Broker Incoming

Port	Purpose
TCP 443	Used to receive activation notifications from Splunk SOAR Cloud via the gRPC tunnel. This is a registration process from the Automation Broker to Spacebridge and provides secure transport to Automation Broker. gRPC is a end to end encryption via TLS 1.2 from Cloud SOAR to Automation Broker

Automation Broker Outbound

Port	Purpose
TCP 443	Used by the automation broker to communicate to Cloud SOAR for action results data. This does not pass through Splunk spacebridge

Standalone Instance

Port	Purpose
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments:

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

S0E - Single Tenant Deployment using Software as a Service with Splunk managed infrastructure with Splunk Integrations with on premise integration

Intentionally left blank for diagram below

SpVA: S0E

Automation Broker Sizing			
	Min	4 Cores	8 GBs
	Recommended	8 Cores	16 GBs

SOAR Cloud region	Cloud Gateway service region
us-west-2	us-east-1
us-east-1	us-east-1
ca-central-1	us-east-1
eu-west-1	eu-central-1
eu-west-2	eu-central-1
eu-west-3	eu-central-1
ap-southeast-1	ap-southeast-2
ap-southeast-2	ap-southeast-2
ap-northeast-1	ap-southeast-2
ap-northeast-2	ap-southeast-2

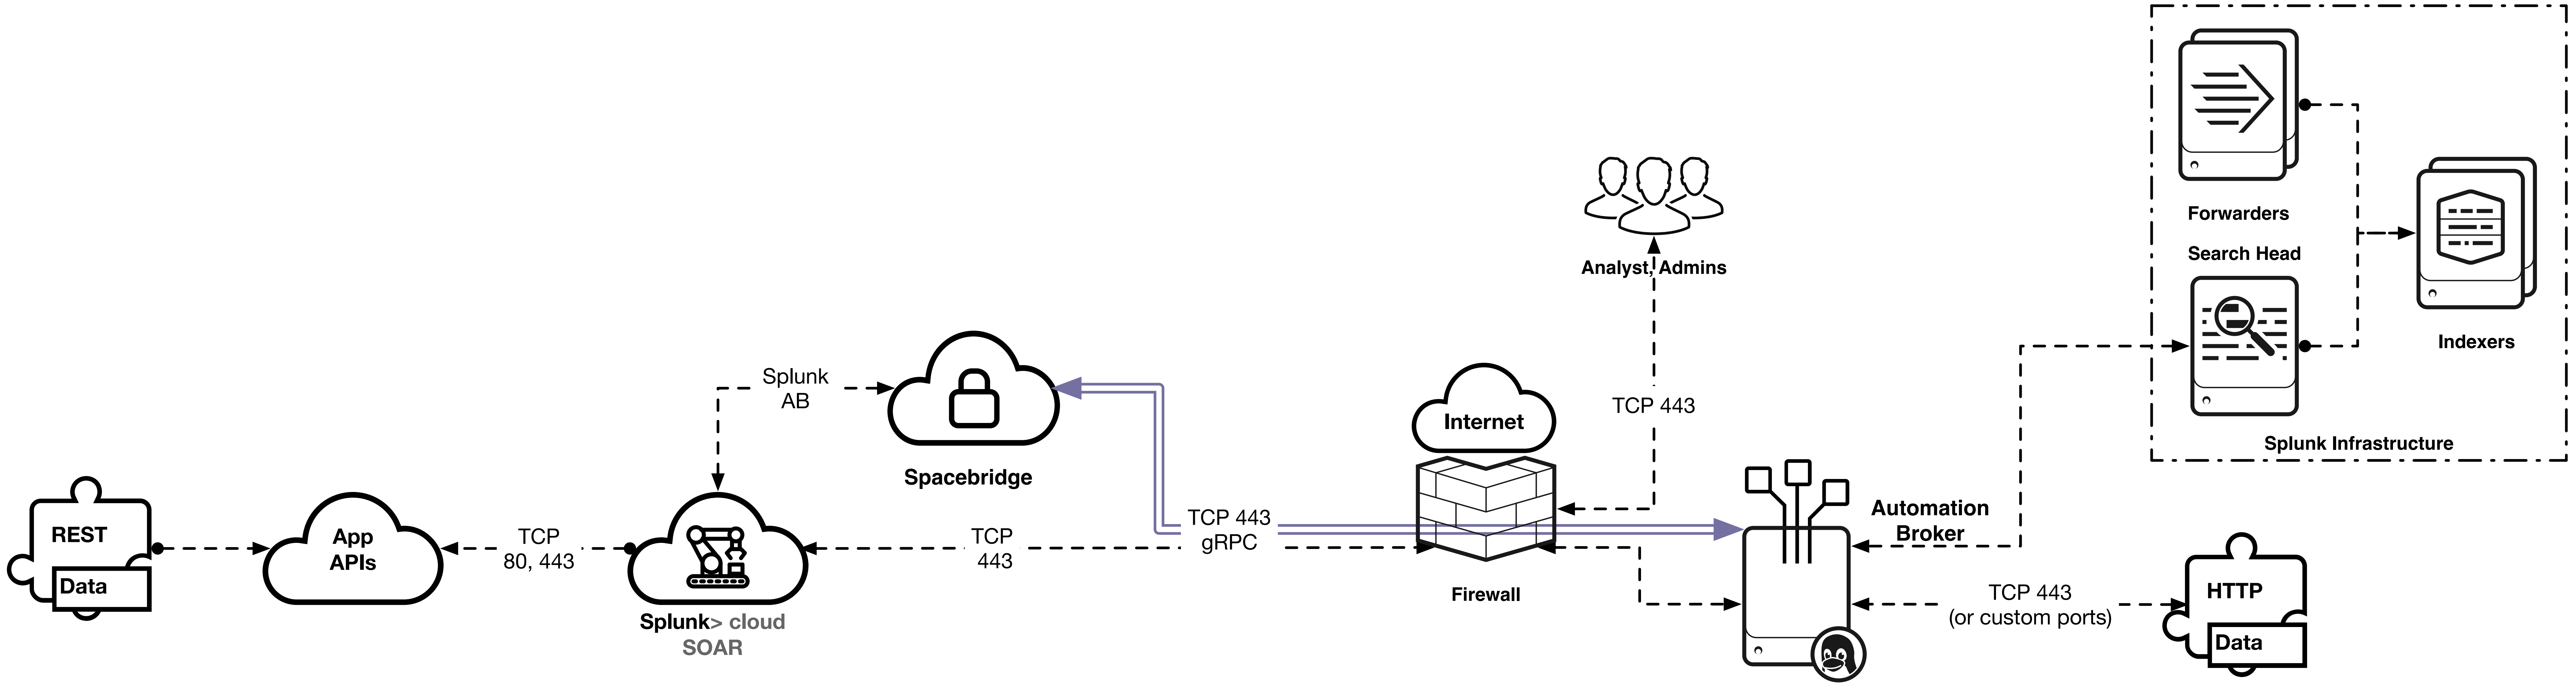

PROS:
Least Administrative overhead
Infrastructure and upgrades managed for customer

CONS:
Performance constrained by number of users

Automation Broker Incoming

Port	Purpose
TCP 443	Used to receive activation notifications from Splunk SOAR Cloud via the gRPC tunnel. This is a registration process from the Automation Broker to Spacebridge and provides secure transport to Automation Broker. gRPC is a end to end encryption via TLS 1.2 from Cloud SOAR to Automation Broker

Automation Broker Outbound

Port	Purpose
TCP 443	Used by the automation broker to communicate to Cloud SOAR for action results data. This does not pass through Splunk spacebridge

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port	Purpose
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments:

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

S0E+ - Single Tenant Deployment using Software as a Service with Splunk managed infrastructure with Splunk Integrations either bring your own cloud

Intentionally left blank for diagram below

SpVA: S0E+

splunk> cloud soar

Automation Broker Sizing	Min	4 Cores	8 GBs
	Recommended	8 Cores	16 GBs

SOAR Cloud region	Cloud Gateway service region
us-west-2	us-east-1
us-east-1	us-east-1
ca-central-1	us-east-1
eu-west-1	eu-central-1
eu-west-2	eu-central-1
eu-west-3	eu-central-1
ap-southeast-1	ap-southeast-2
ap-southeast-2	ap-southeast-2
ap-northeast-1	ap-southeast-2
ap-northeast-2	ap-southeast-2

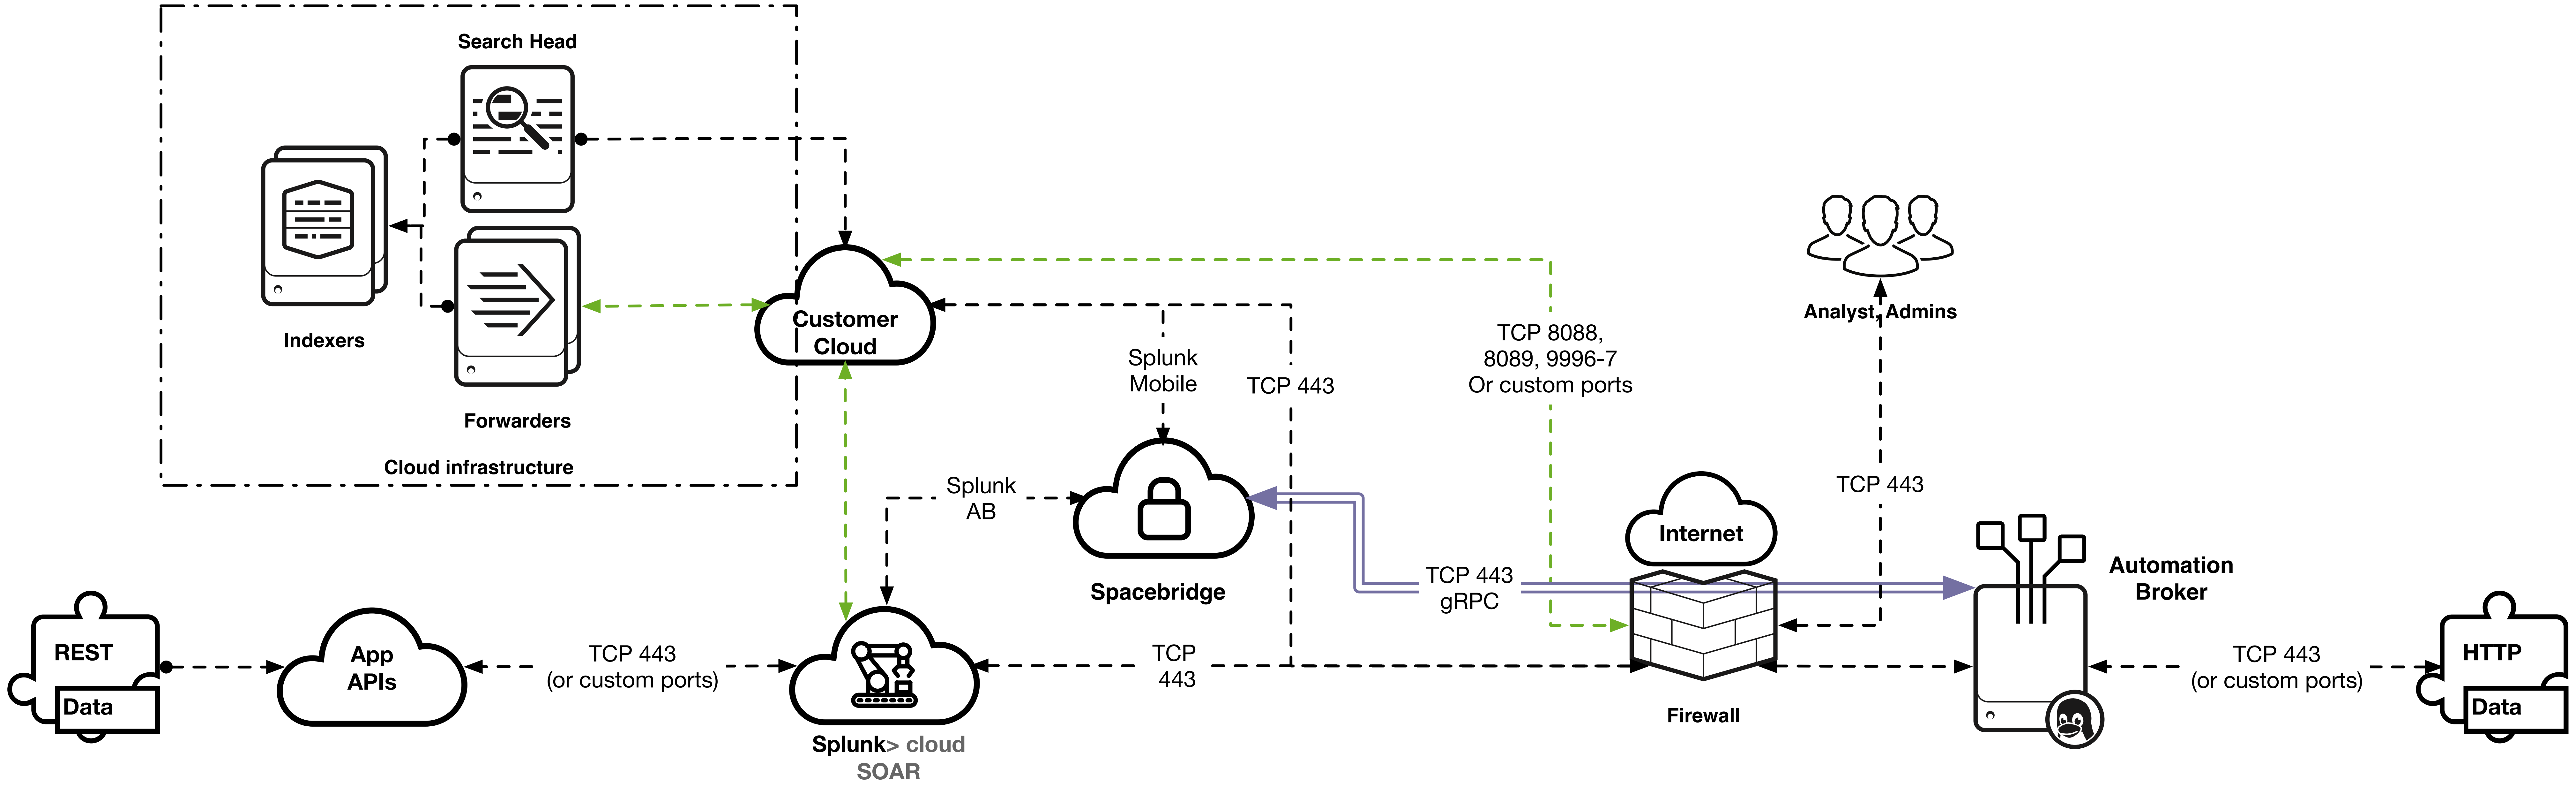

PROS:
Least Administrative overhead
Infrastructure and upgrades managed for customer

CONS:
Performance constrained by number of users

Automation Broker Incoming

Port	Purpose
TCP 443	Used to receive activation notifications from Splunk SOAR Cloud via the gRPC tunnel. This is a registration process from the Automation Broker to Spacebridge and provides secure transport to Automation Broker. gRPC is a end to end encryption via TLS 1.2 from Cloud SOAR to Automation Broker

Automation Broker Outbound

Port	Purpose
TCP 443	Used by the automation broker to communicate to Cloud SOAR for action results data. This does not pass through Splunk spacebridge

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port	Purpose
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments:

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

S0CE - Single Tenant Deployment using Software as a Service with Splunk managed infrastructure and Splunk Cloud Integrations

Intentionally left blank for diagram below

SpVA: S0CE

splunk> cloud soar

Automation Broker Sizing			
	Min	4 Cores	8 GBs
	Recommended	8 Cores	16 GBs

SOAR Cloud region	Cloud Gateway service region
us-west-2	us-east-1
us-east-1	us-east-1
ca-central-1	us-east-1
eu-west-1	eu-central-1
eu-west-2	eu-central-1
eu-west-3	eu-central-1
ap-southeast-1	ap-southeast-2
ap-southeast-2	ap-southeast-2
ap-northeast-1	ap-southeast-2
ap-northeast-2	ap-southeast-2

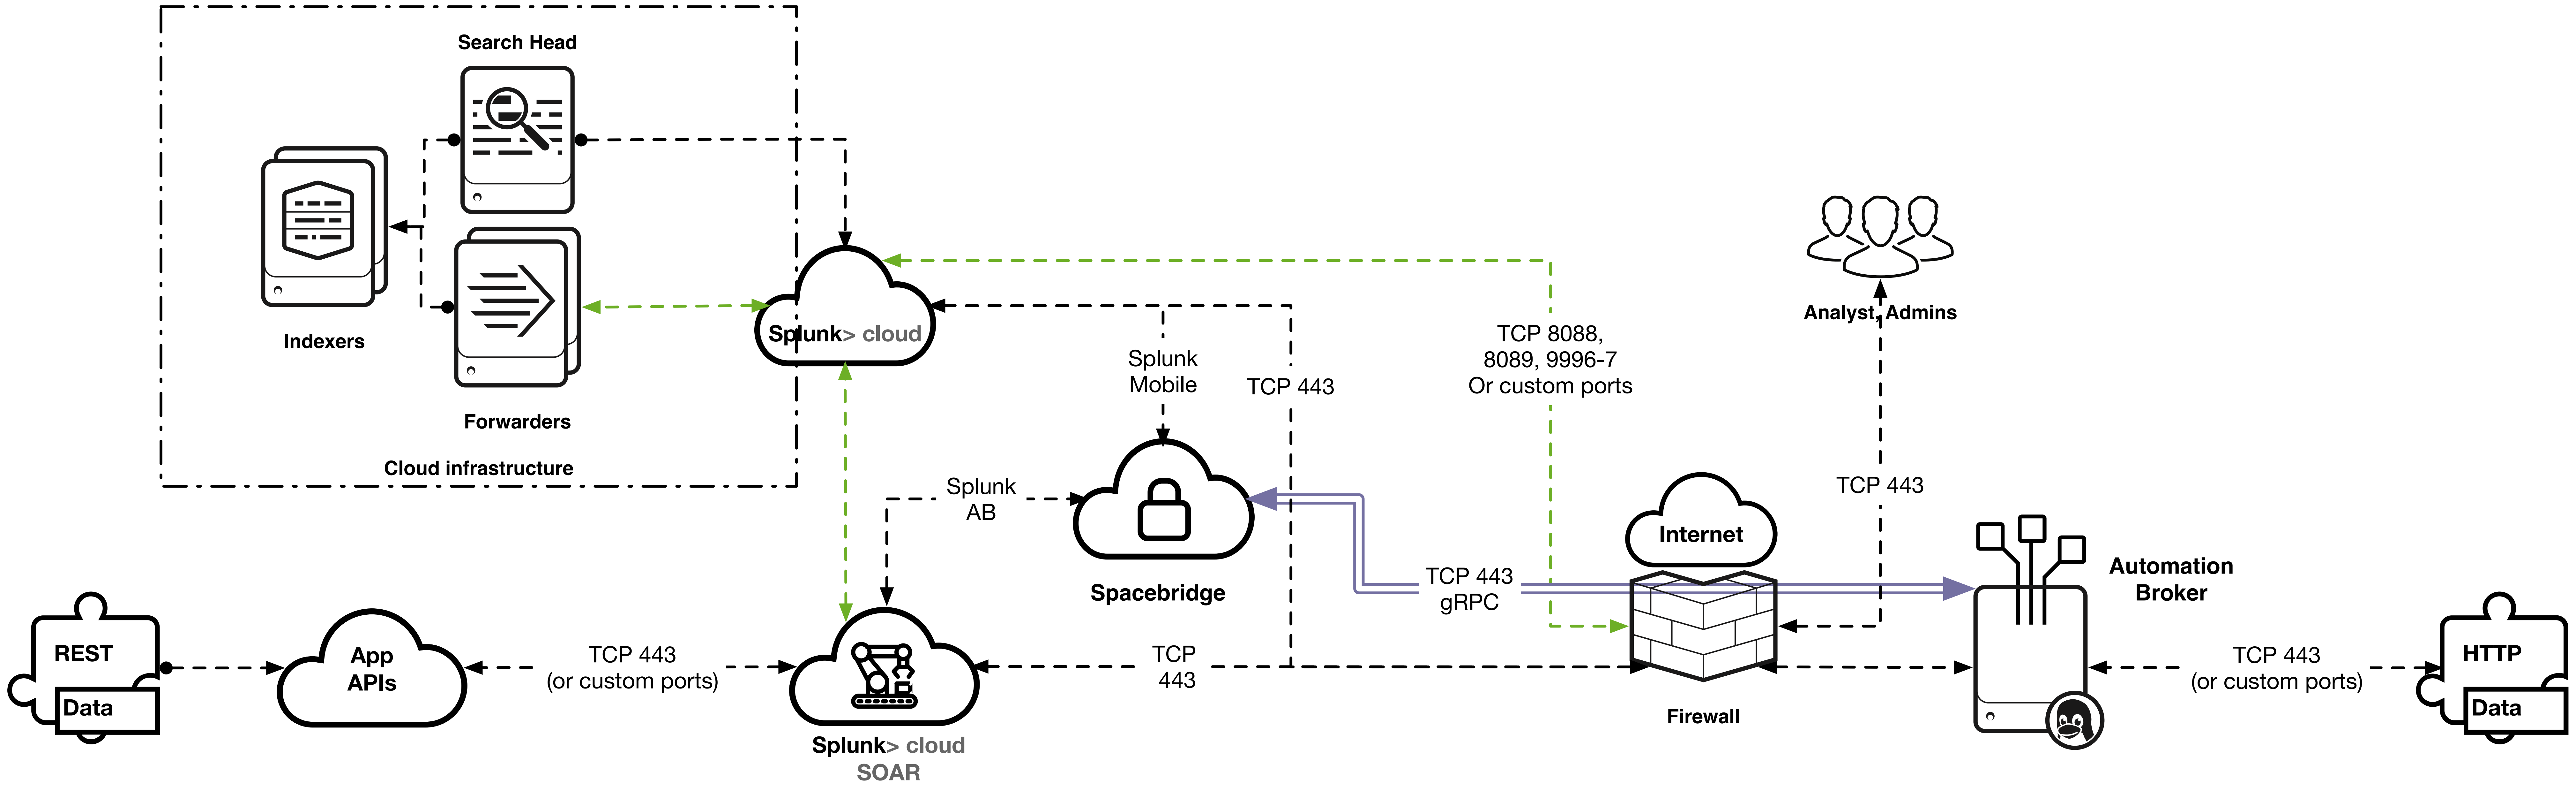

Automation Broker Incoming

Port — Purpose

TCP 443 — Used to receive activation notifications from Splunk SOAR Cloud via the gRPC tunnel. This is a registration process from the Automation Broker to Spacebridge and provides secure transport to Automation Broker. gRPC is a end to end encryption via TLS 1.2 from Cloud SOAR to Automation Broker

PROS:
Least Administrative overhead
Infrastructure and upgrades managed for customer

CONS:
Performance constrained by number of users

Automation Broker Outbound

Port — Purpose

TCP 443 — Used by the automation broker to communicate to Cloud SOAR for action results data. This does not pass through Splunk spacebridge

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port — Purpose

TCP 443 — HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments:

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

S1 - Single Server Deployment using embedded Splunk Integration

Intentionally left blank for diagram below

SpVA: S1

splunk> phantom

Splunk Phantom Stand-Alone Sizing Chart

Instance Type	Workloads with active playbooks	CPU cores	Memory GB
Large	>7000 events per hour	32	64
Medium	Up to 7000 events per hour	16	32
Small ***Development ONLY***	Up to 4000 events per hour	8	16
Tiny ***Development ONLY***	< 4000 events per hour	8	8

Recommended Sizing

Default OVA Build	**Based on size (GB):**	**200**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-var	/	40
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	45
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	45
/dev/mapper/centos-tmp	/tmp	10
/dev/mapper/centos-var	/var	15
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	1
/dev/mapper/centos-var_tmp	/var/tmp	15
/dev/mapper/centos-home	/home	20

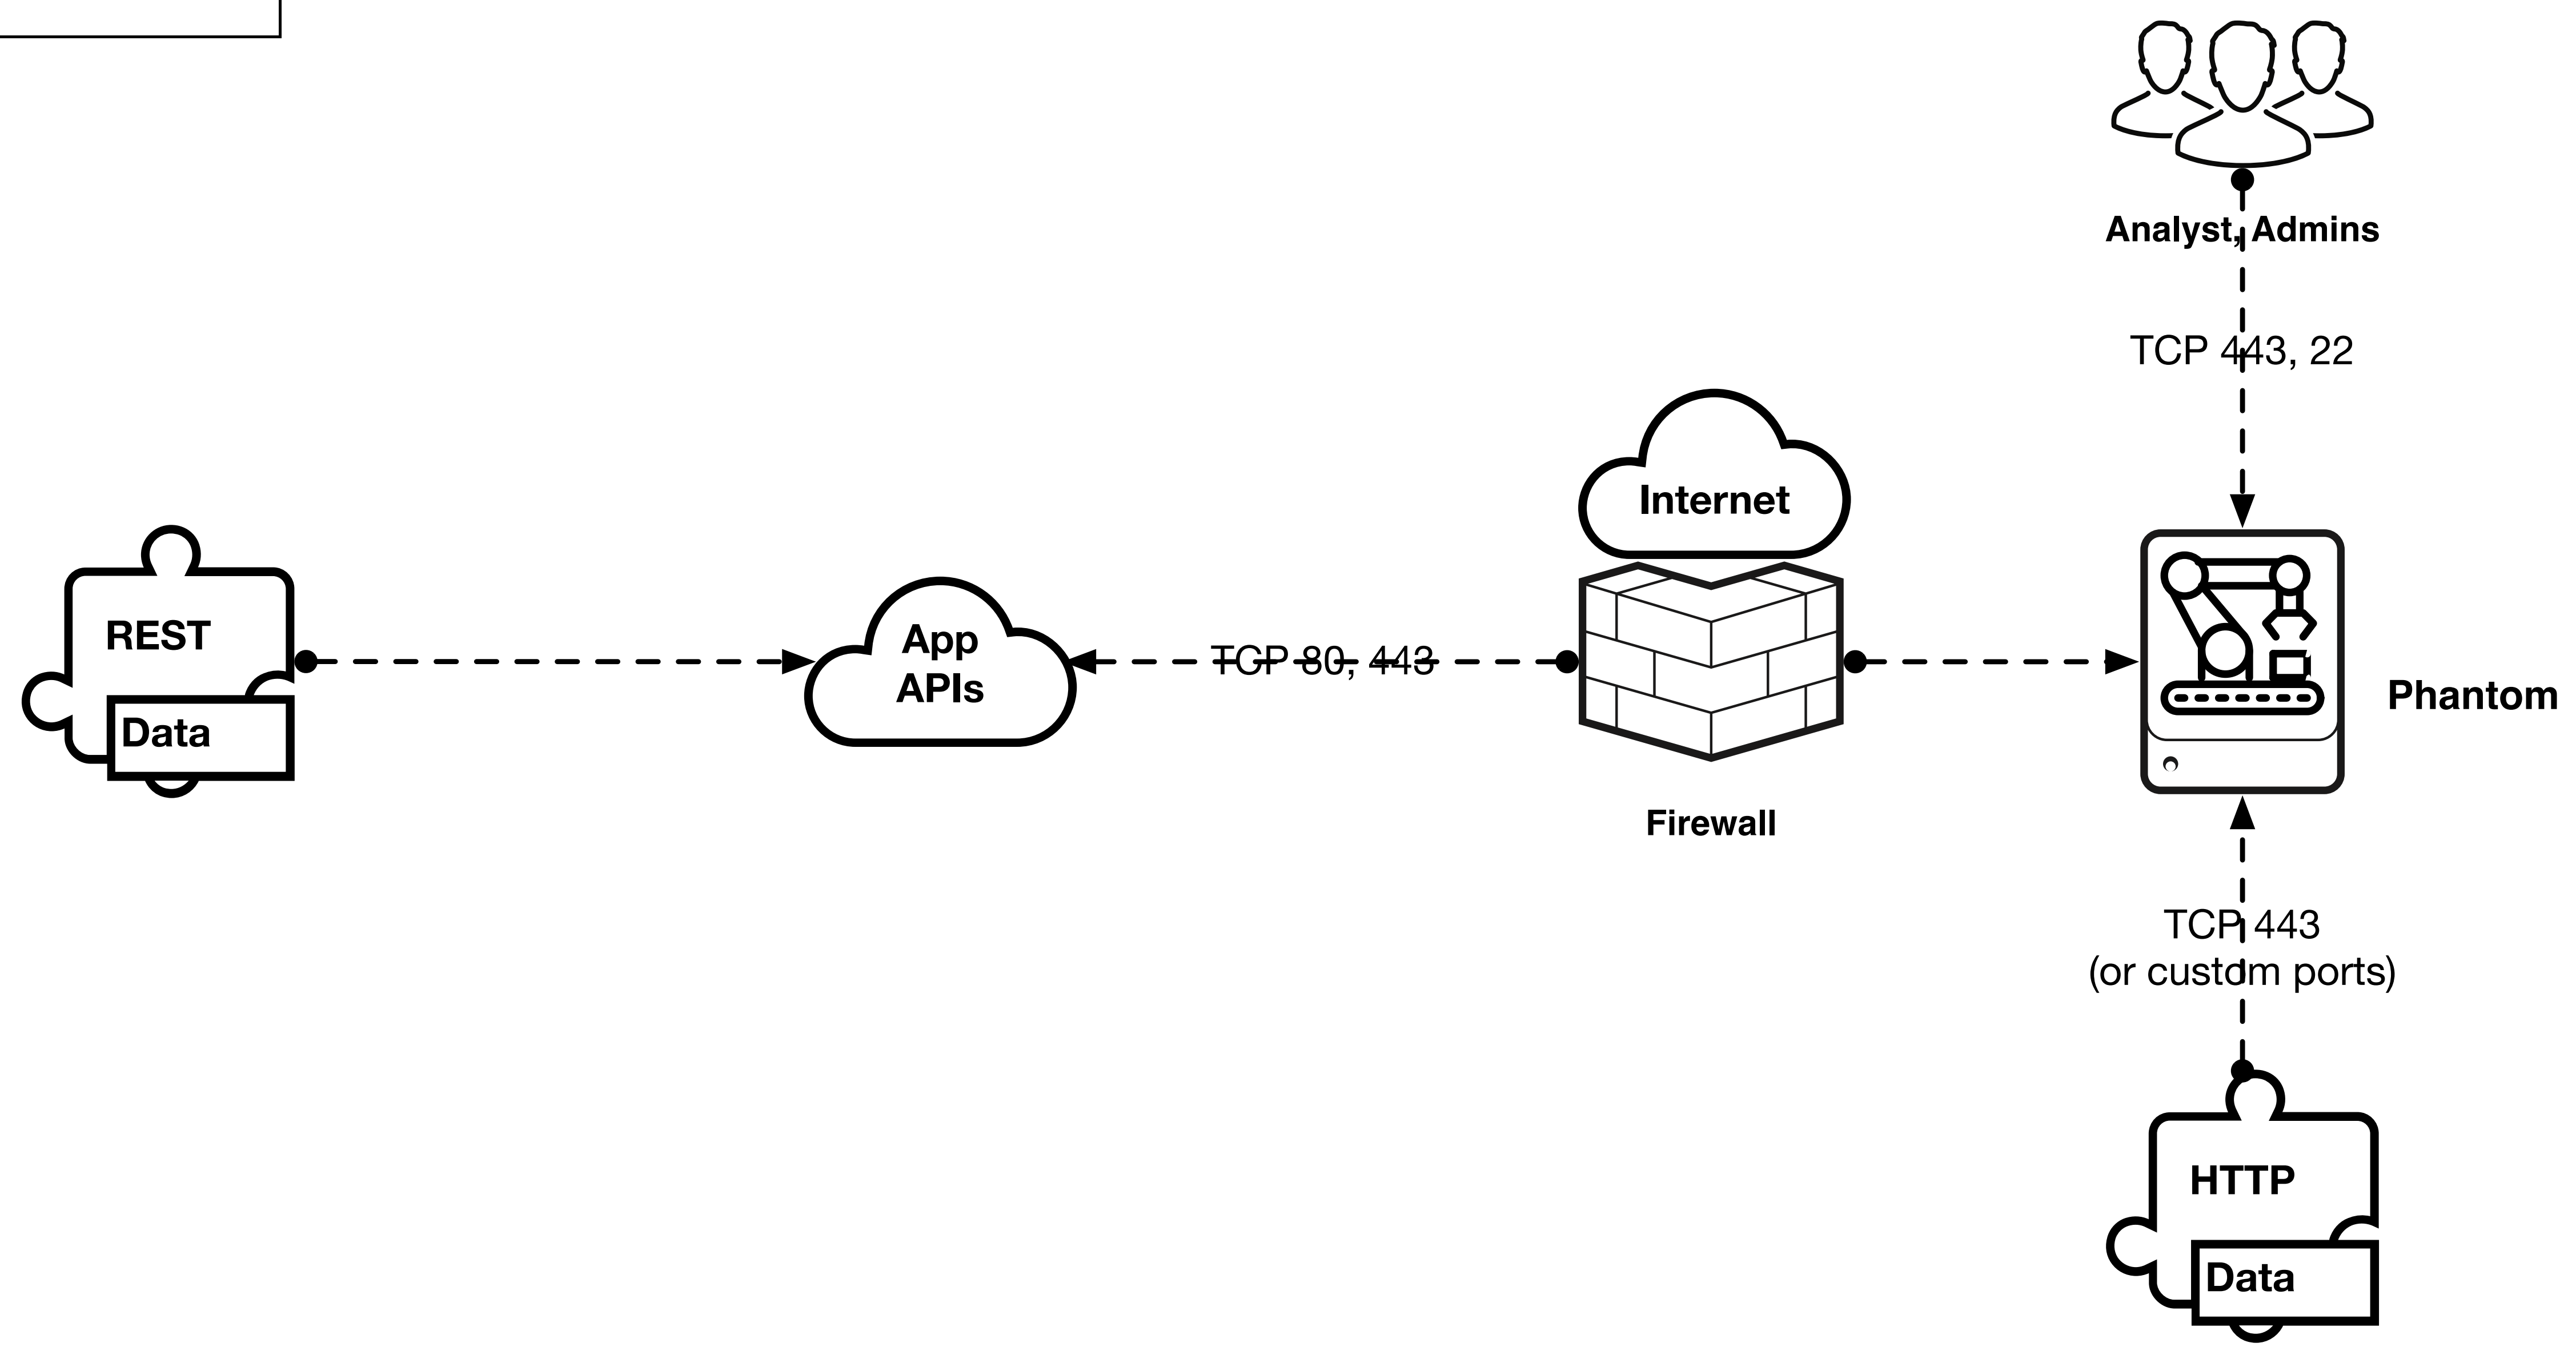

PROS:
Least Administrative overhead
Least TCO going towards production

CONS:
Performance constrained by number of available connections
Backup and Recover is customer responsibility
No failover mechanisms

Comments:

REVISIONS			
REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

DRAWN BY Architect

ISSUED TO Approver

Company Name

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

S1E - Single Server Deployment using Splunk Enterprise Integration

Intentionally left blank for diagram below

SpVA: S1E

splunk> phantom

Splunk Phantom Stand-Alone Sizing Chart

Instance Type	Workloads with active playbooks	CPU cores	Memory GB
Large	>7000 events per hour	32	64
Medium	Up to 7000 events per hour	16	32
Small ***Development ONLY***	Up to 4000 events per hour	8	16
Tiny ***Development ONLY***	< 4000 events per hour	8	8

Recommended Sizing

Production Build Recommend Drive Mappings	Based on size (GB):	1024
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	204.8
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	245.8
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	307.2
/dev/mapper/centos-tmp	/tmp	51.2
/dev/mapper/centos-var	/var	51.2
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	5.1
/dev/mapper/centos-var_tmp	/var/tmp	51.2
/dev/mapper/centos-home	/home	102.4
	Total:	**1018.9**

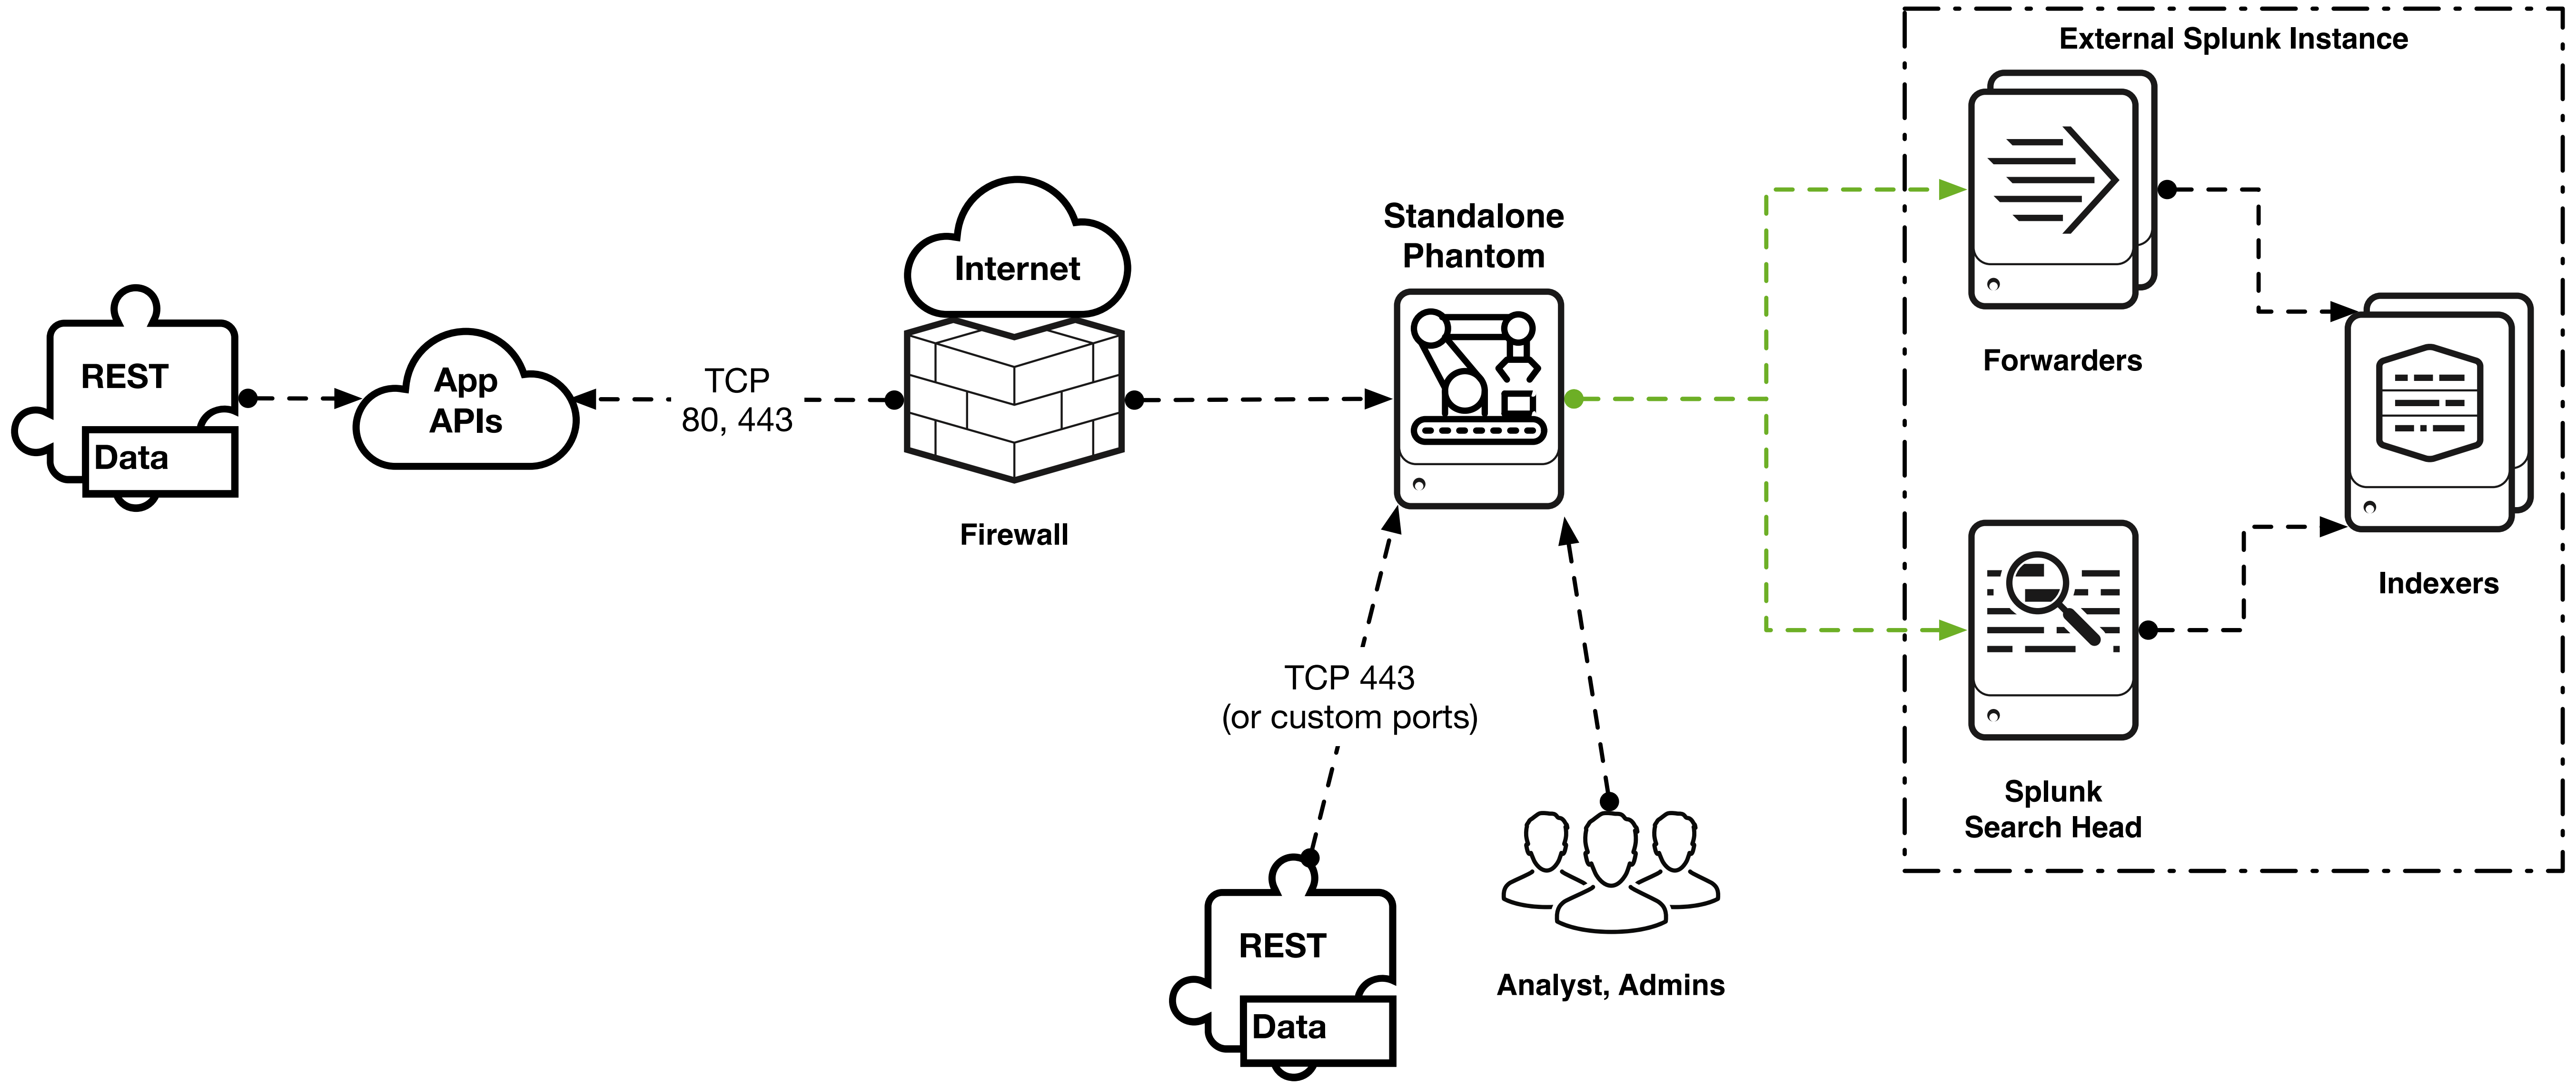

PROS:
Least Administrative overhead with improved UI performance
Least number of resource contention
Easiest to manage and gradual expansion increase process

CONS:
Performance constrained by number of available connections
No failover mechanisms

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation for Single Instance deployments

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build capable for most loads < 1500 events per day and for teams using Case Management and Headless usage	6 Aug 18	

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

X1 - Single Server Deployment using external Shared Services with embedded Splunk

Intentionally left blank for diagram below

SpVA: X1

splunk> phantom

Splunk Phantom Stand-Alone Sizing Chart

Instance Type	Workloads with active playbooks	CPU cores	Memory GB
Large	>7000 events per hour	32	64
Medium	Up to 7000 events per hour	16	32
Small	Up to 4000 events per hour	8	16
Tiny	< 4000 events per hour	8	8

Recommended Sizing

Default OVA Build	Based on size (GB):	200
Device	Mountpoint	Size (GB)
/dev/mapper/centos-var	/	40
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	45
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	45
/dev/mapper/centos-tmp	/tmp	10
/dev/mapper/centos-var	/var	15
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	1
/dev/mapper/centos-var_tmp	/var/tmp	15
/dev/mapper/centos-home	/home	20

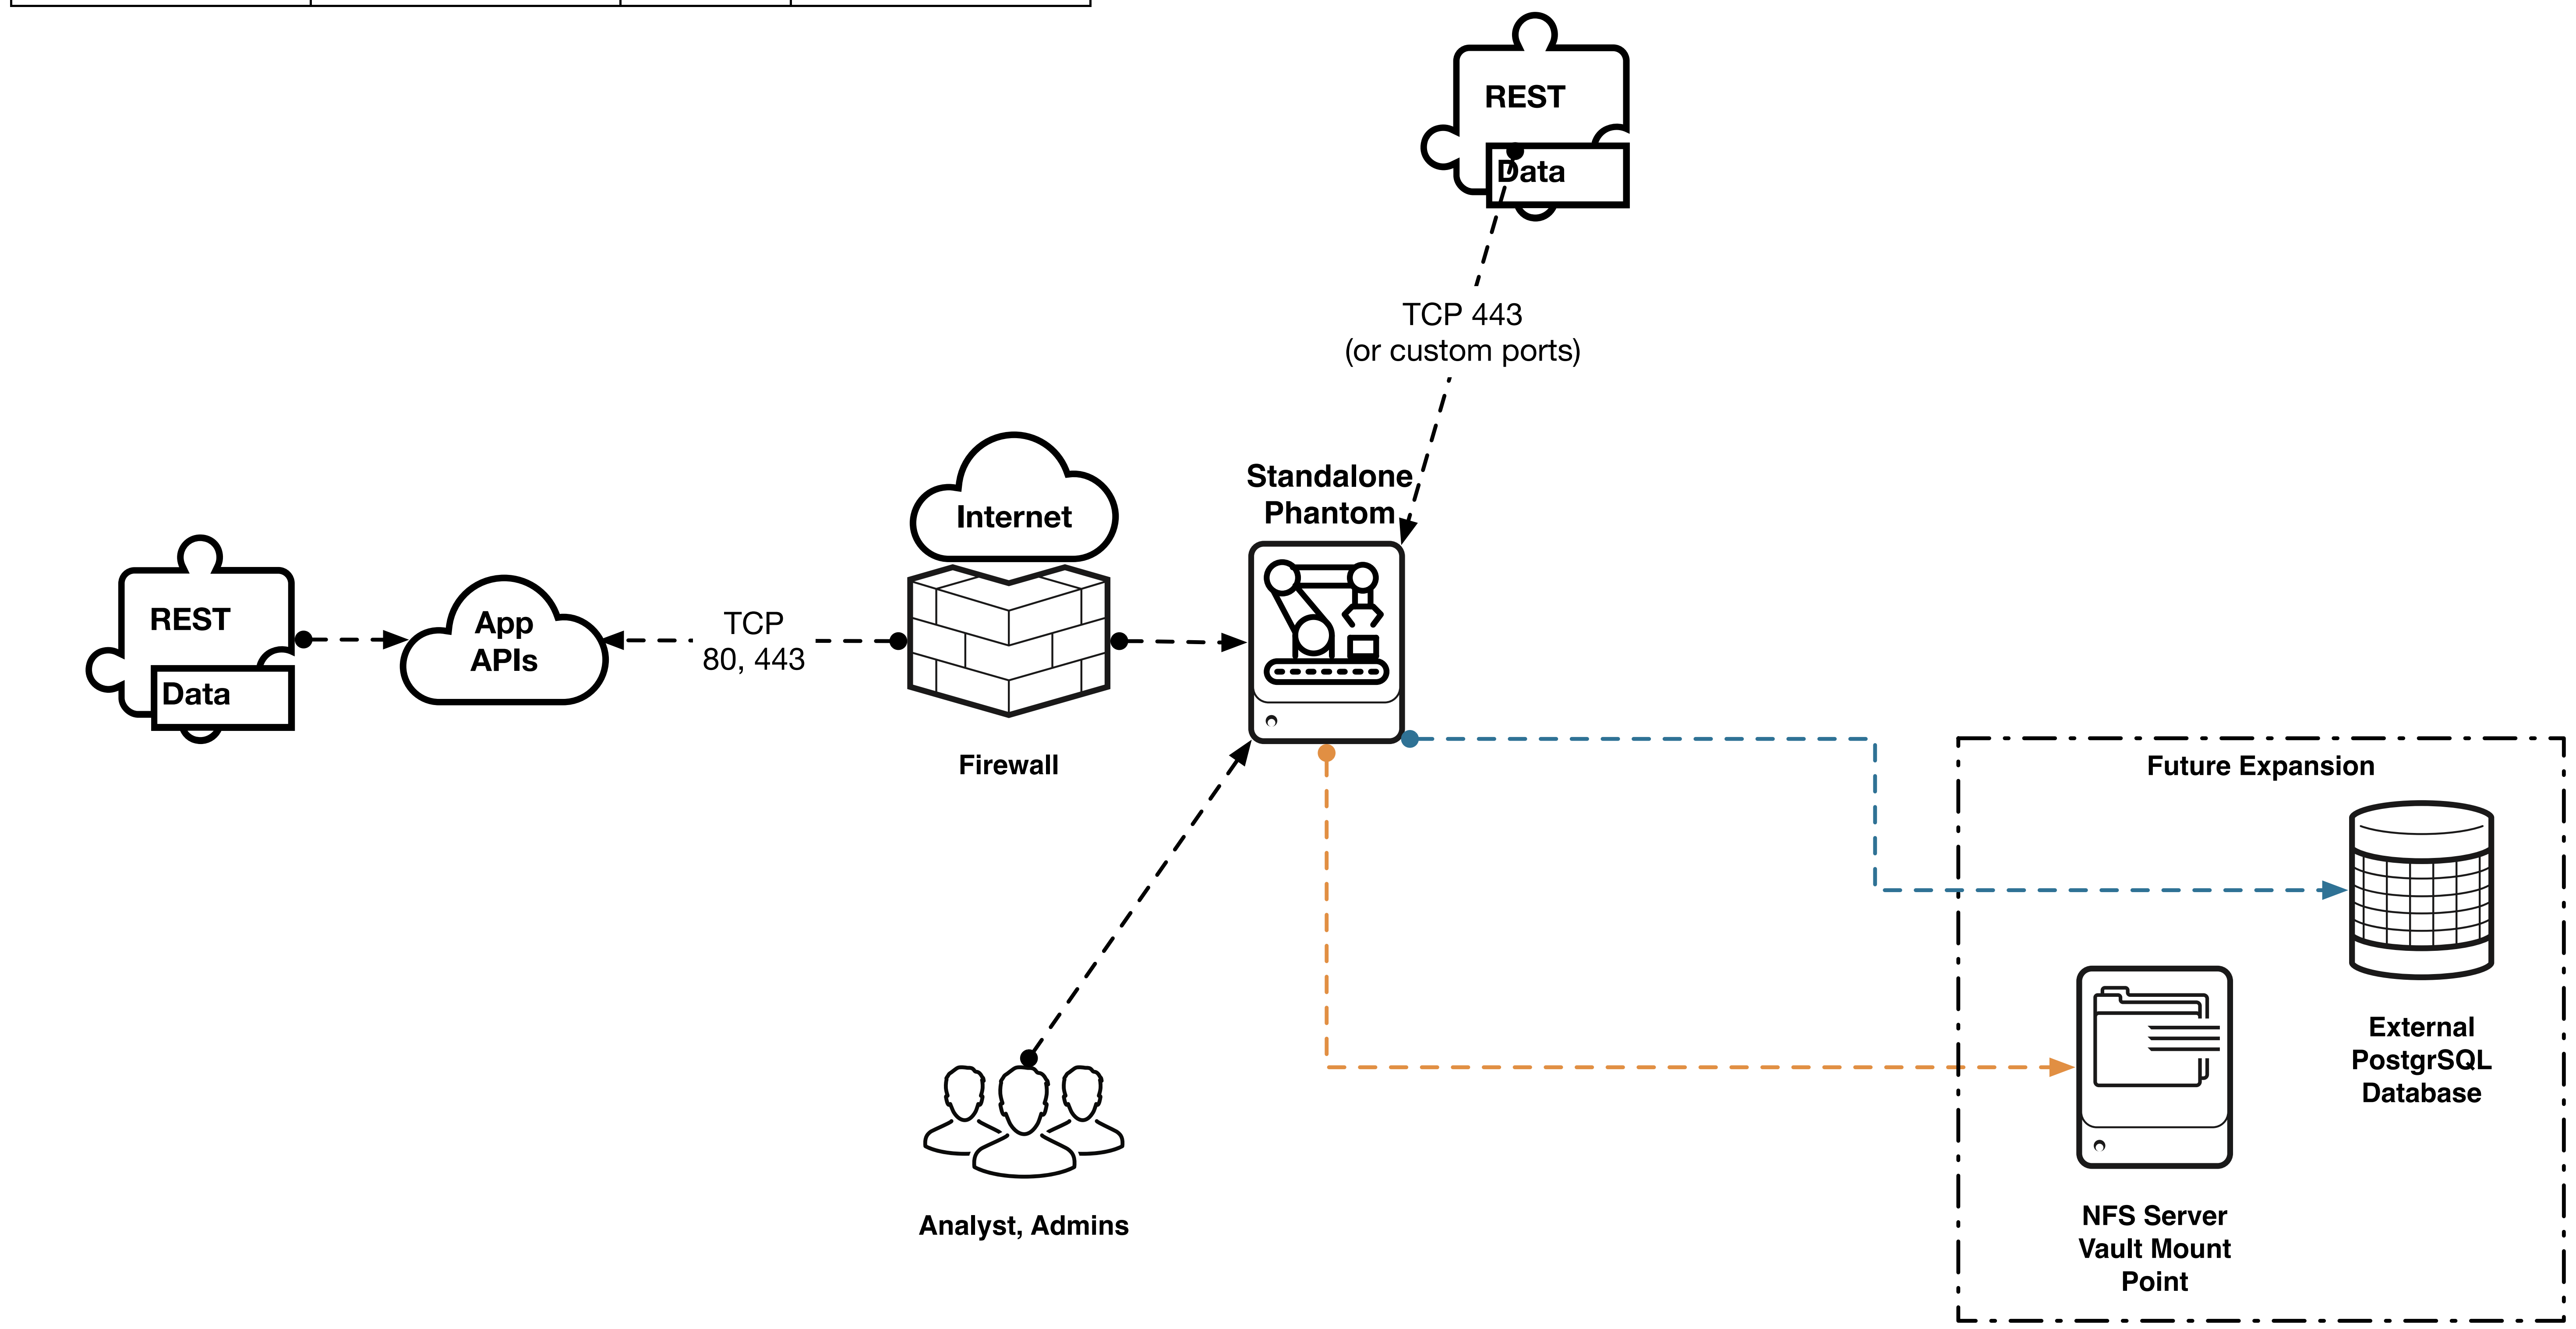

PROS:
Least Administrative overhead with improved UI performance
Least number of resource contention
Easiest to manage and gradual expansion increase process

CONS:
Performance constrained by number of available connections
No failover mechanisms

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

POSTGRES Instance

Port	Purpose
TCP 22	Used for administering the OS that POSTGRES is running on
TCP 5432	PostgreSQL Service. Can be blocked if the DB server is a different host than the shared services node.

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This model is only used for prior to clustering and single instance support without warm standby.

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build capable for most loads < 1500 events per day and for teams using Case Management and Headless usage	6 Aug 18	

X1E - Single Server Deployment using external Shared Services with Splunk Enterprise Integration

Intentionally left blank for diagram below

Splunk Phantom Stand-Alone Sizing Chart

Instance Type	Workloads with active playbooks	CPU cores	Memory GB
Large	>7000 events per hour	32	64
Medium	Up to 7000 events per hour	16	32
Small	Up to 4000 events per hour	8	16
Tiny	< 4000 events per hour	8	8

Recommended Sizing

SpVA: X1E

Default OVA Build	Based on size (GB):	200
Device	Mountpoint	Size (GB)
/dev/mapper/centos-var	/	40
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	45
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	45
/dev/mapper/centos-tmp	/tmp	10
/dev/mapper/centos-var	/var	15
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	1
/dev/mapper/centos-var_tmp	/var/tmp	15
/dev/mapper/centos-home	/home	20

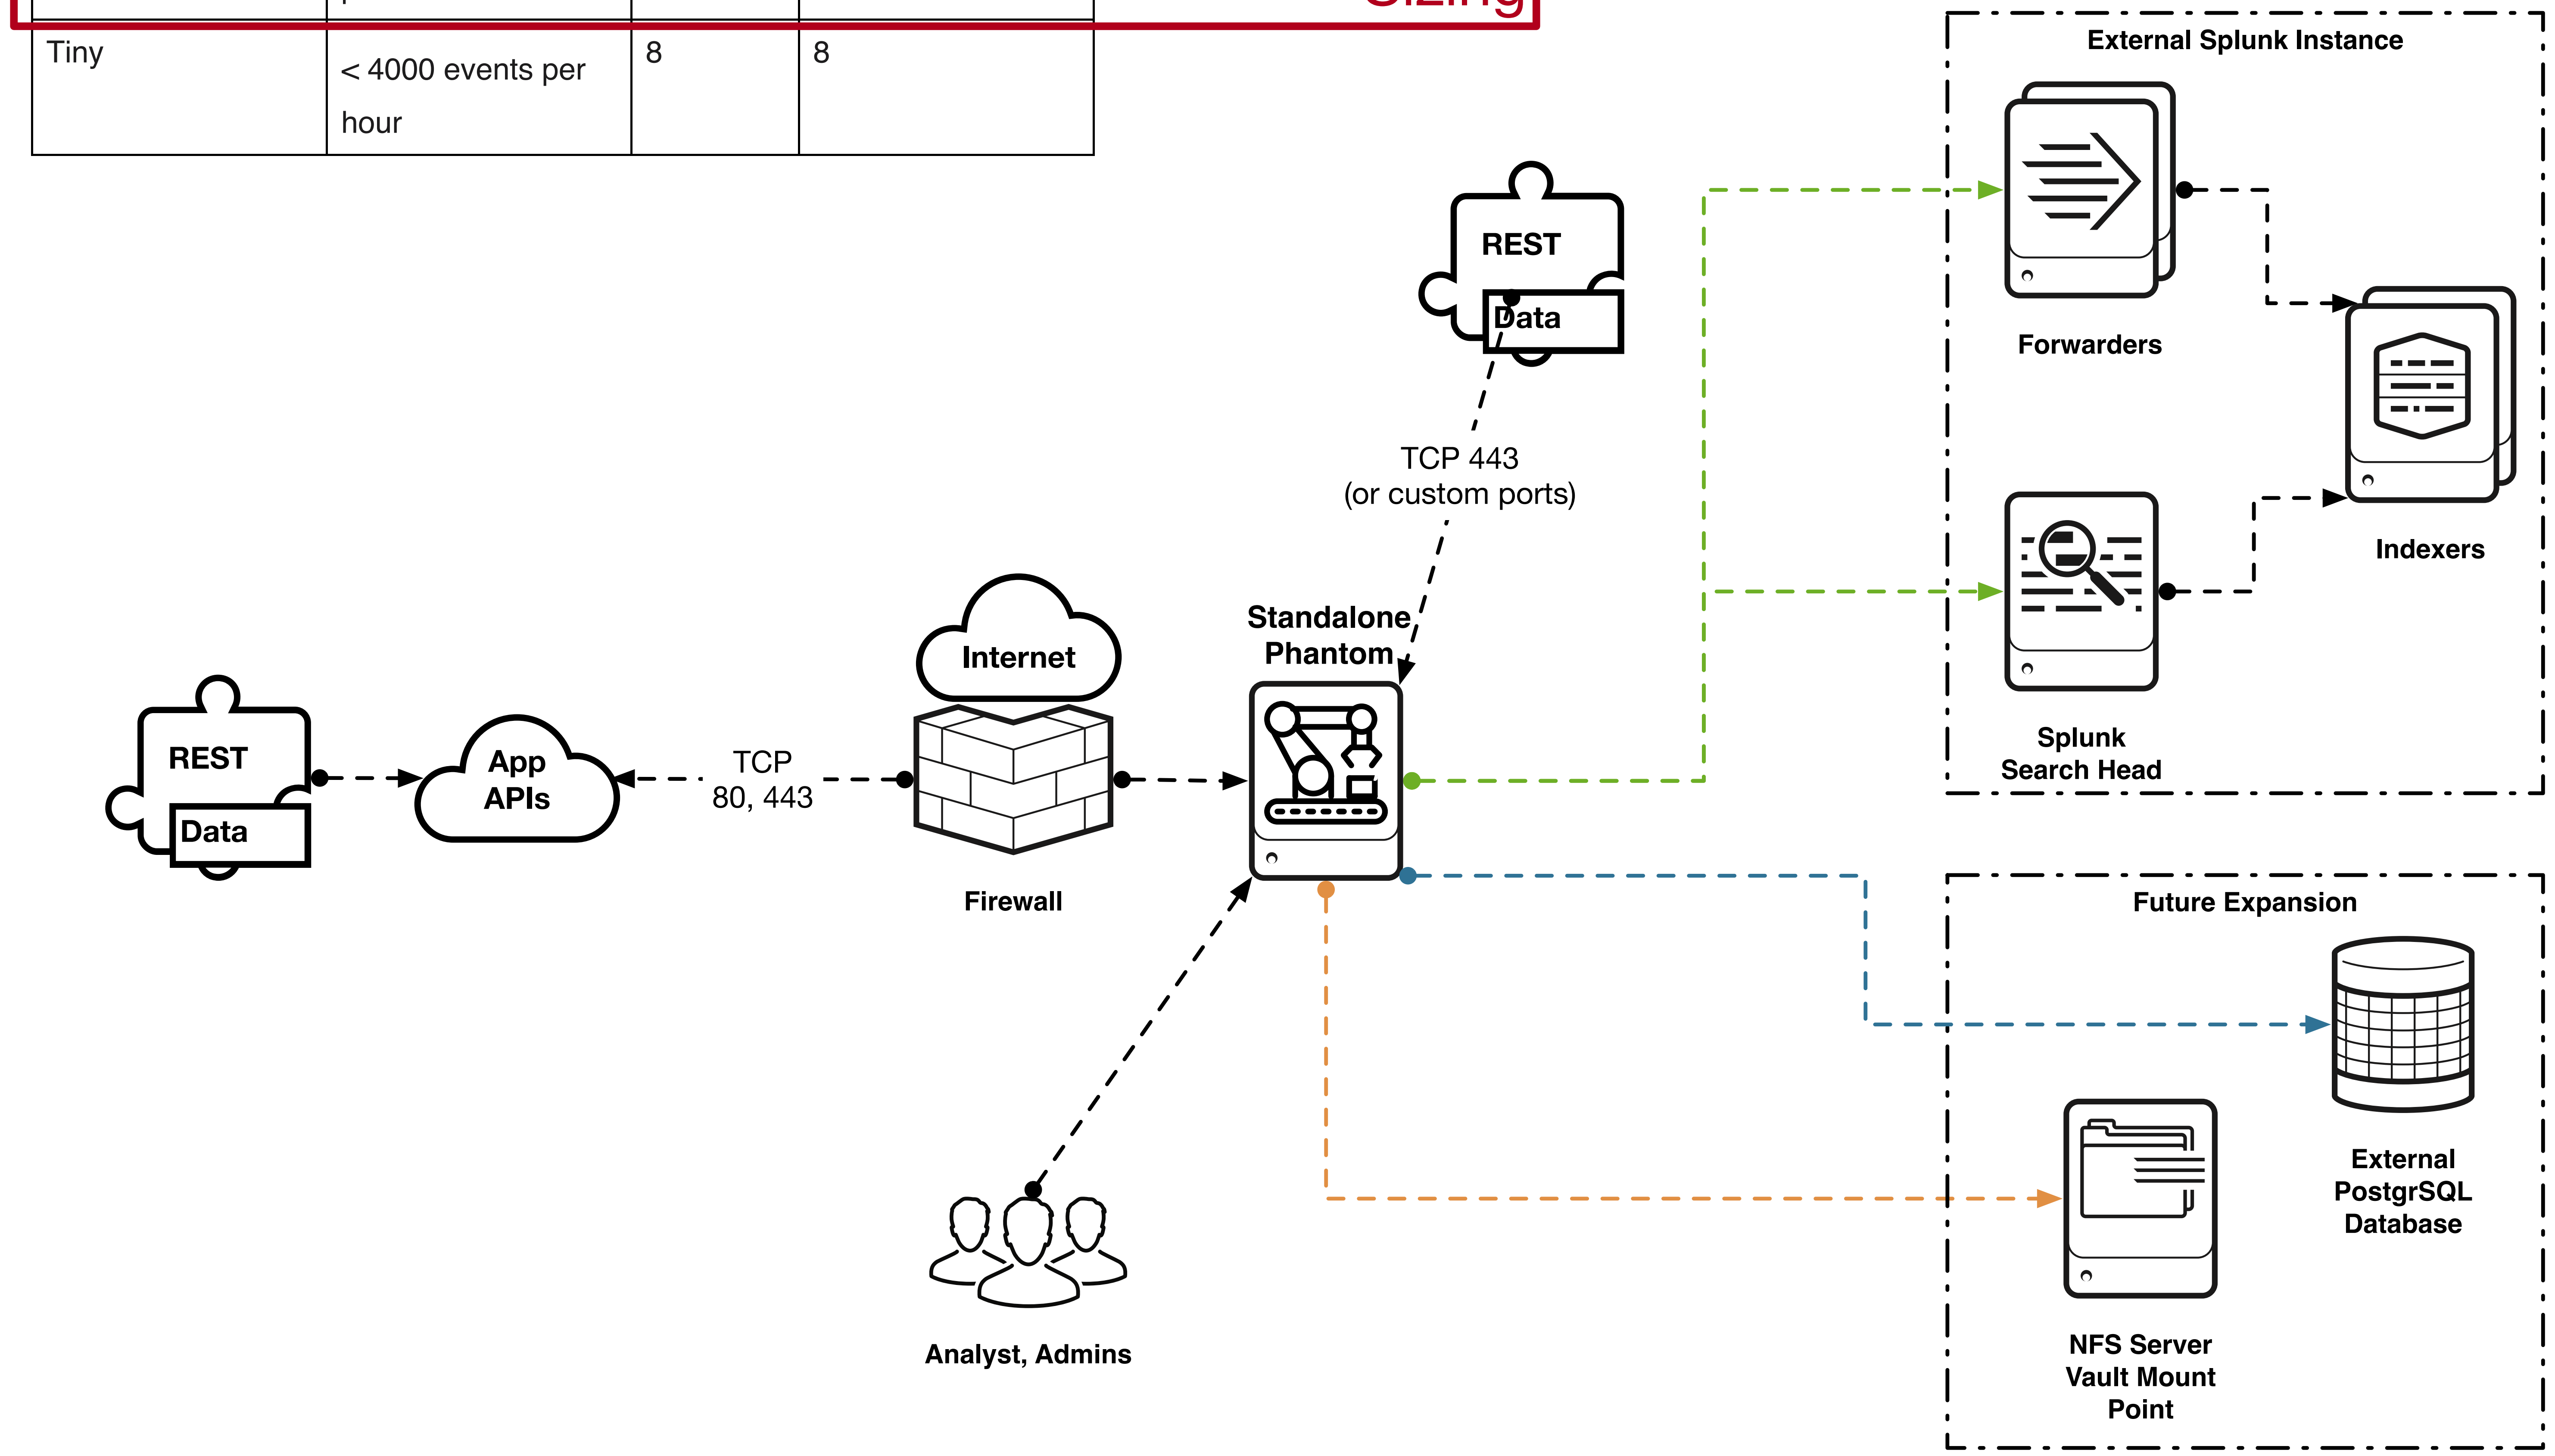

PROS:
Least Administrative overhead with improved UI performance
Least number of resource contention
Easiest to manage and gradual expansion increase process

CONS:
Performance constrained by number of available connections
No failover mechanisms

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

POSTGRES Instance

Port	Purpose
TCP 22	Used for administering the OS that POSTGRES is running on
TCP 5432	PostgreSQL Service. Can be blocked if the DB server is a different host than the shared services node.

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This model is only used for prior to clustering and single instance support without warm standby.

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build capable for most loads < 1500 events per day and for teams using Case Management and Headless usage	6 Aug 18	

D2 - Distributed Warm Standby Deployment using embedded Splunk integration

Intentionally left blank for diagram below

SpVA: D2

splunk> phantom

Splunk Phantom Stand-Alone Sizing Chart

Instance Type	Workloads with active playbooks	CPU cores	Memory GB
Large	>7000 events per hour	32	64
Medium	Up to 7000 events per hour	16	32
Small *Development ONLY*	Up to 4000 events per hour	8	16
Tiny *Development ONLY*	< 4000 events per hour	8	8

Recommended Sizing

Production Build Recommend Drive Mappings	Based on size (GB):	1024
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	204.8
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	245.8
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	307.2
/dev/mapper/centos-tmp	/tmp	51.2
/dev/mapper/centos-var	/var	51.2
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	5.1
/dev/mapper/centos-var_tmp	/var/tmp	51.2
/dev/mapper/centos-home	/home	102.4
	Total:	1018.9

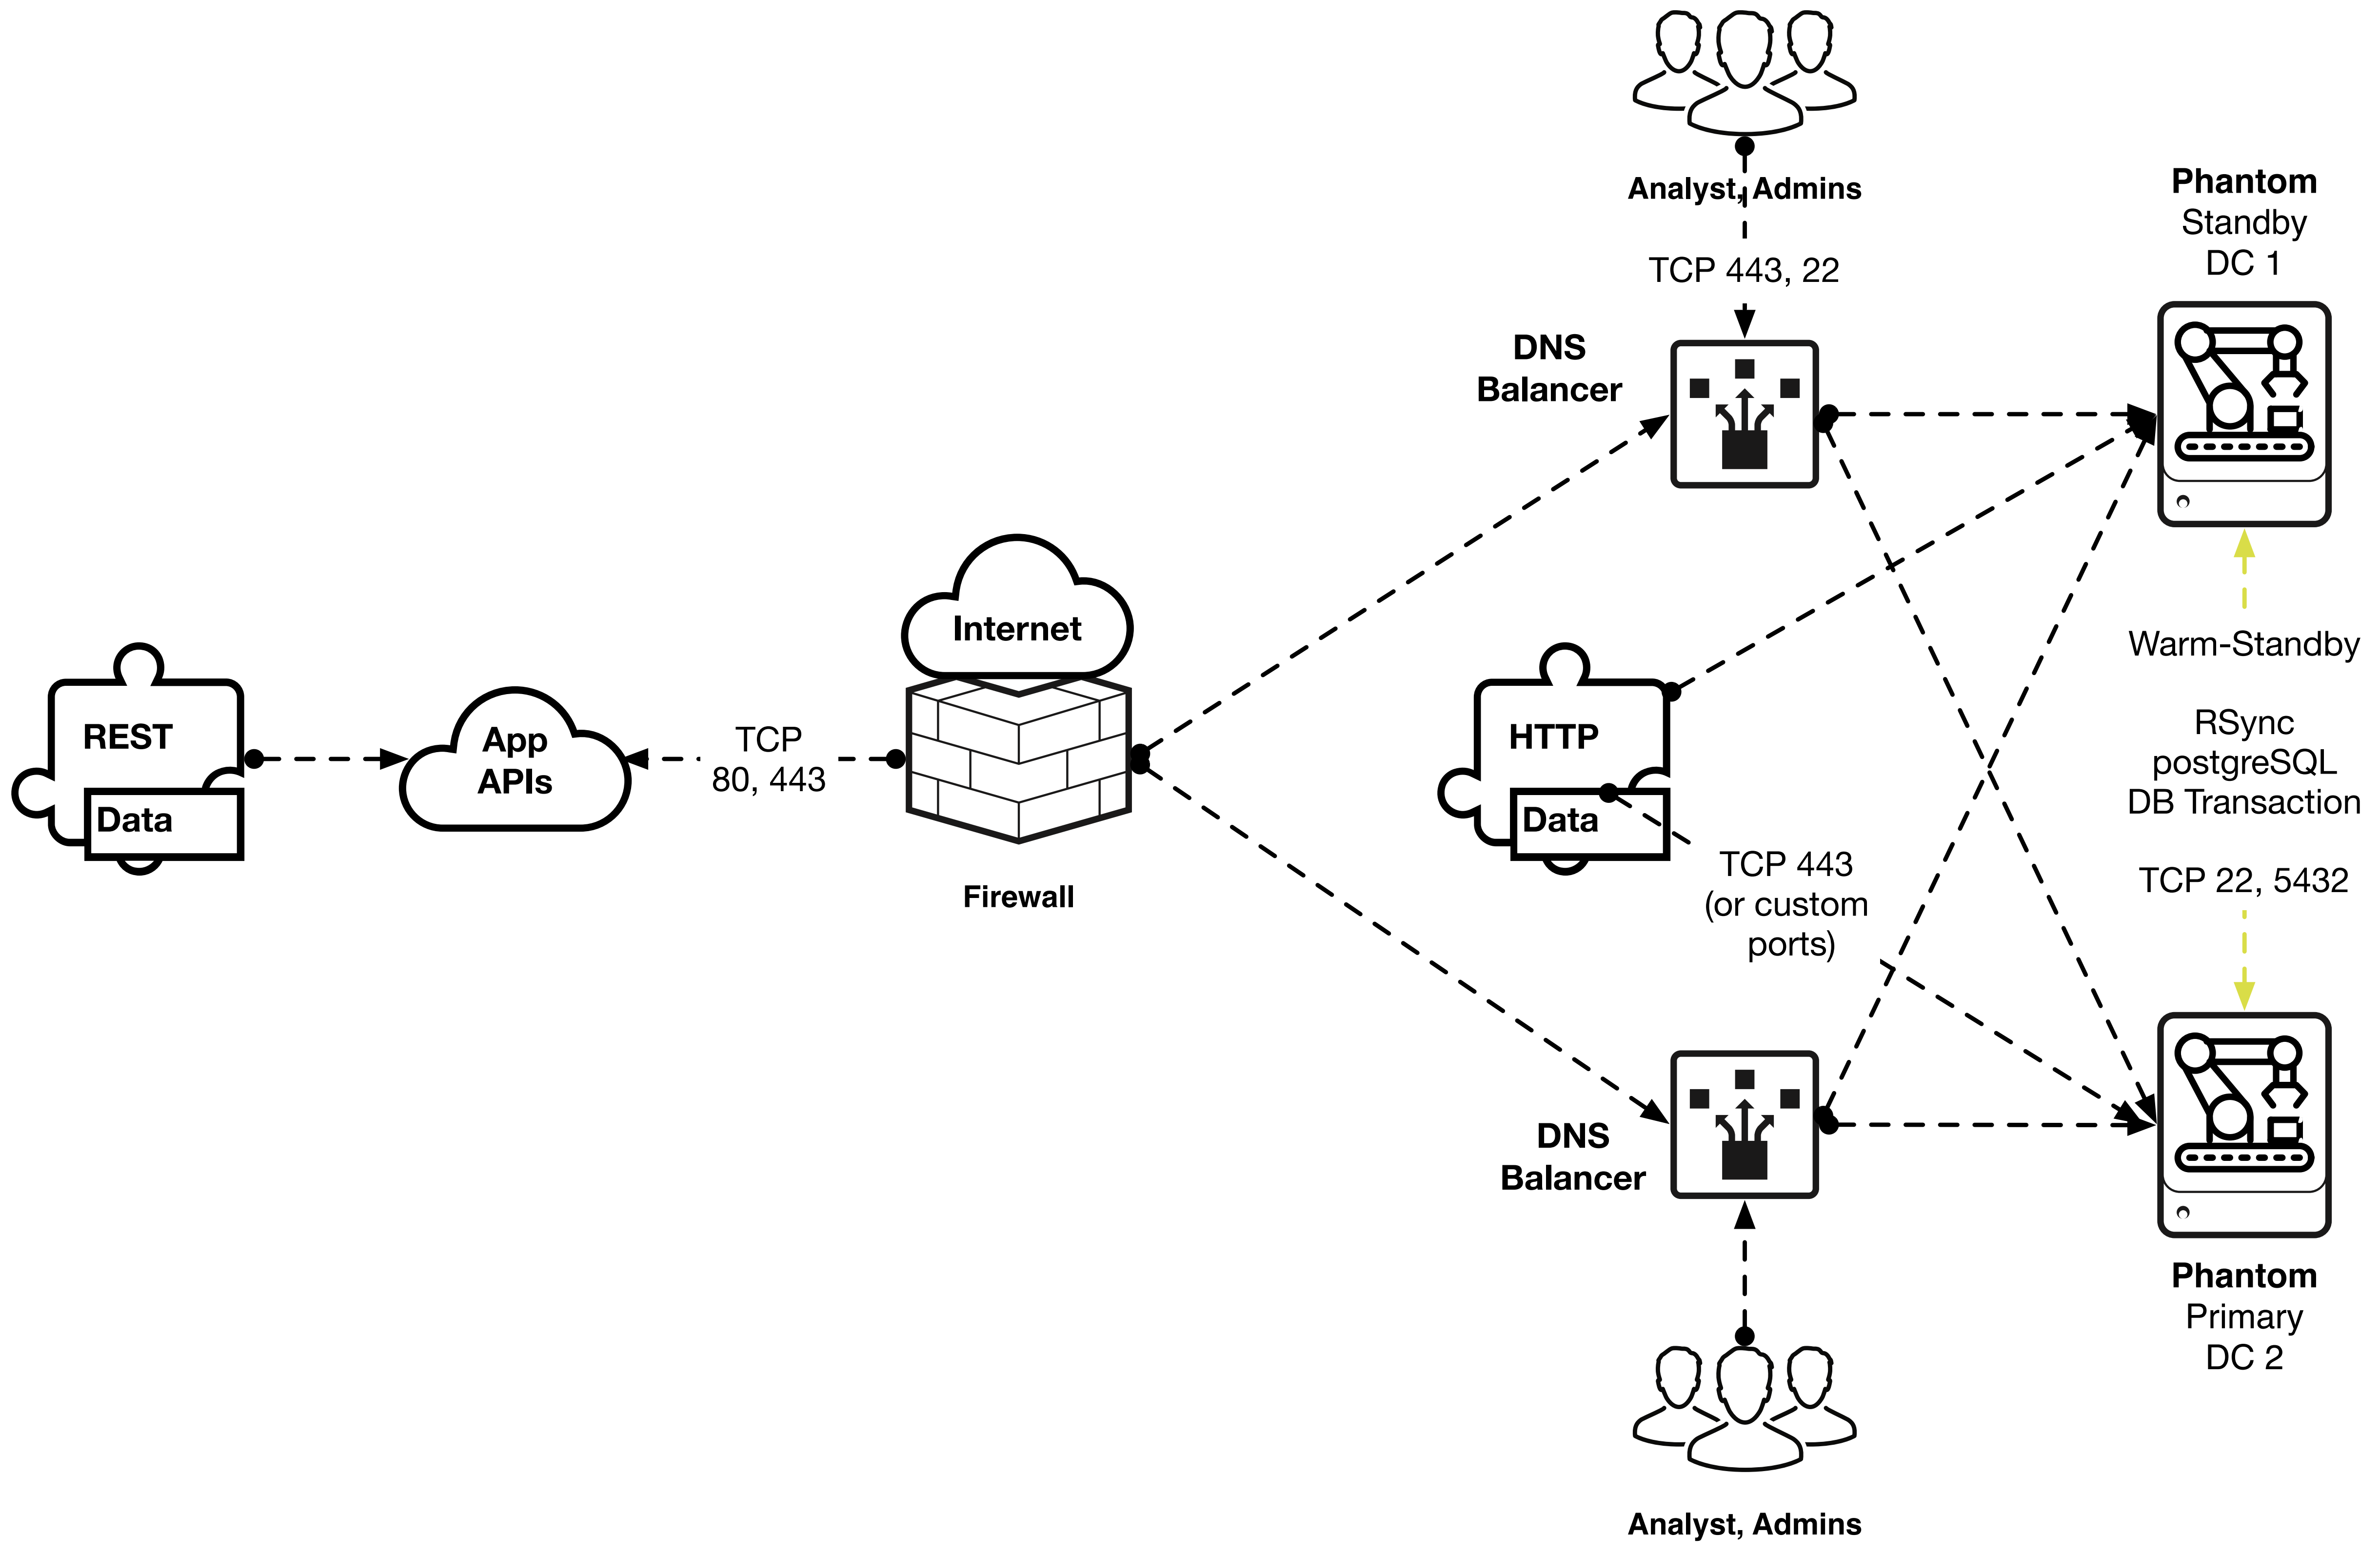

Warm Standby Process

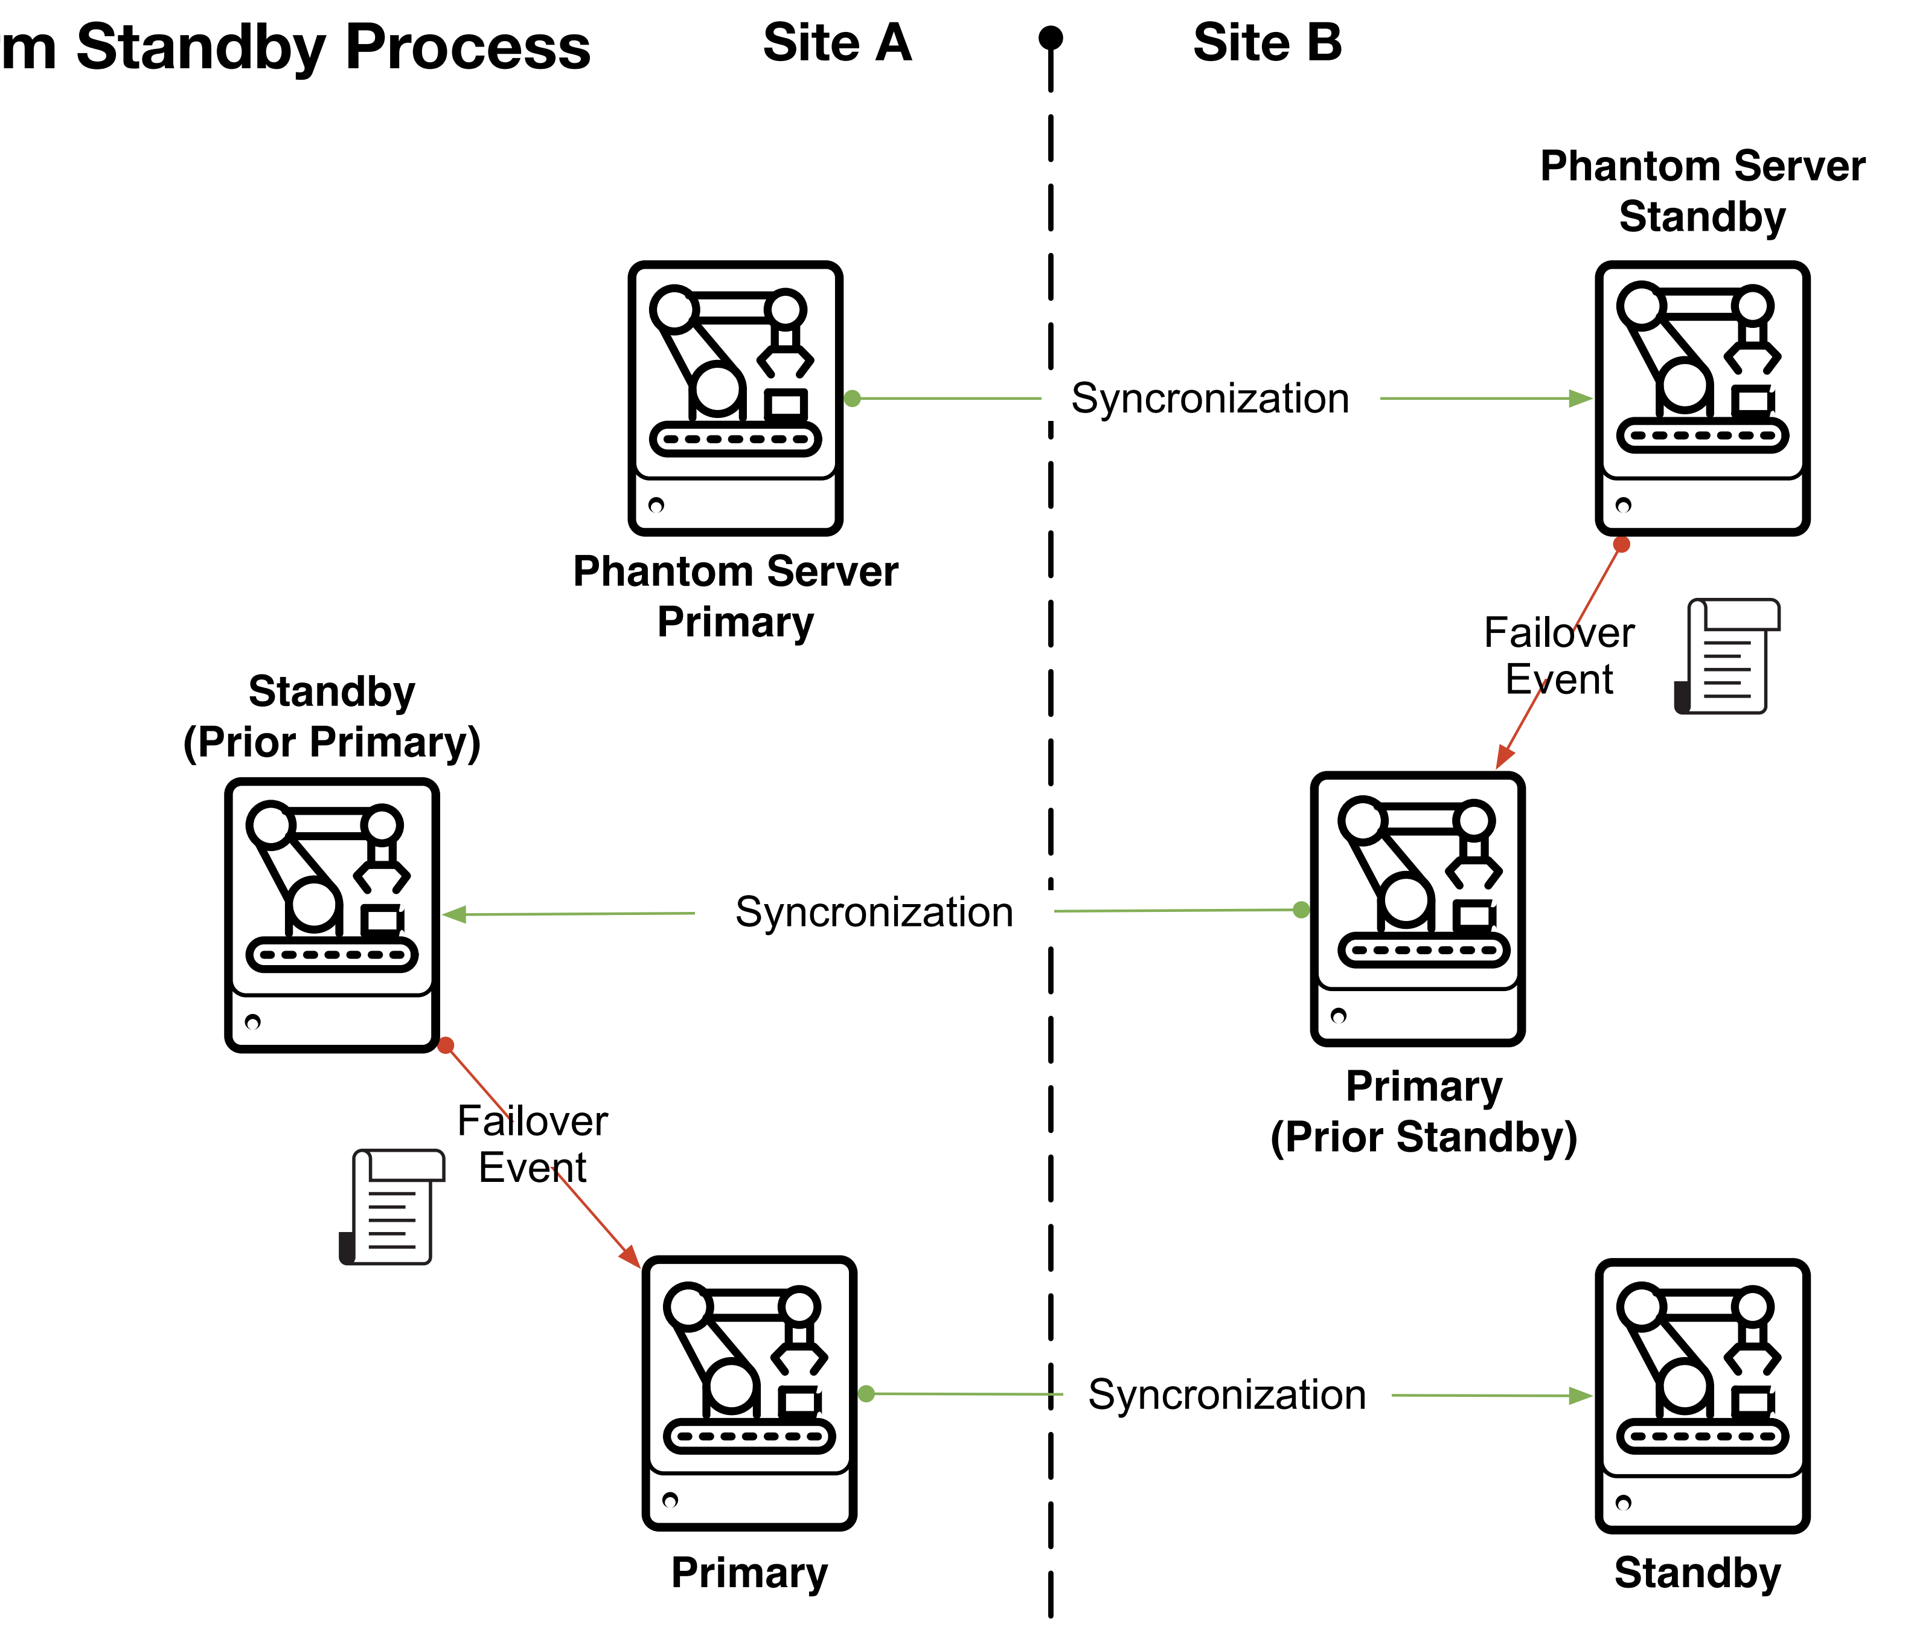

PROS:
Least number of resources
Native multi-site redundancy and high availability
Easiest to manage and gradual expansion increase process

CONS:
Performance constrained by number of available connections
Failover is scripted and can be semi-automated or manual

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation for Multiple Site Survivability

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build capable for most loads < 1500 events per day and for teams using Case Management	3 MAR 19	

DRAWN BY	Architect
ISSUED TO	Customer

Company | Company Confidential

D2E - Distributed Warm Standby Deployment using embedded Splunk with customer provided Infrastructure

Intentionally left blank for diagram below

Splunk Phantom Stand-Alone Sizing Chart

Instance Type	Workloads with active playbooks	CPU cores	Memory GB
Large	>7000 events per hour	32	64
Medium	Up to 7000 events per hour	16	32
Small ***Development ONLY***	Up to 4000 events per hour	8	16
Tiny ***Development ONLY***	< 4000 events per hour	8	8

Recommended Sizing

SpVA: D2E

splunk> phantom

Production Build Recommend Drive Mappings	Based on size (GB):	1024
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	204.8
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	245.8
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	307.2
/dev/mapper/centos-tmp	/tmp	51.2
/dev/mapper/centos-var	/var	51.2
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	5.1
/dev/mapper/centos-var_tmp	/var/tmp	51.2
/dev/mapper/centos-home	/home	102.4
	Total:	1018.9

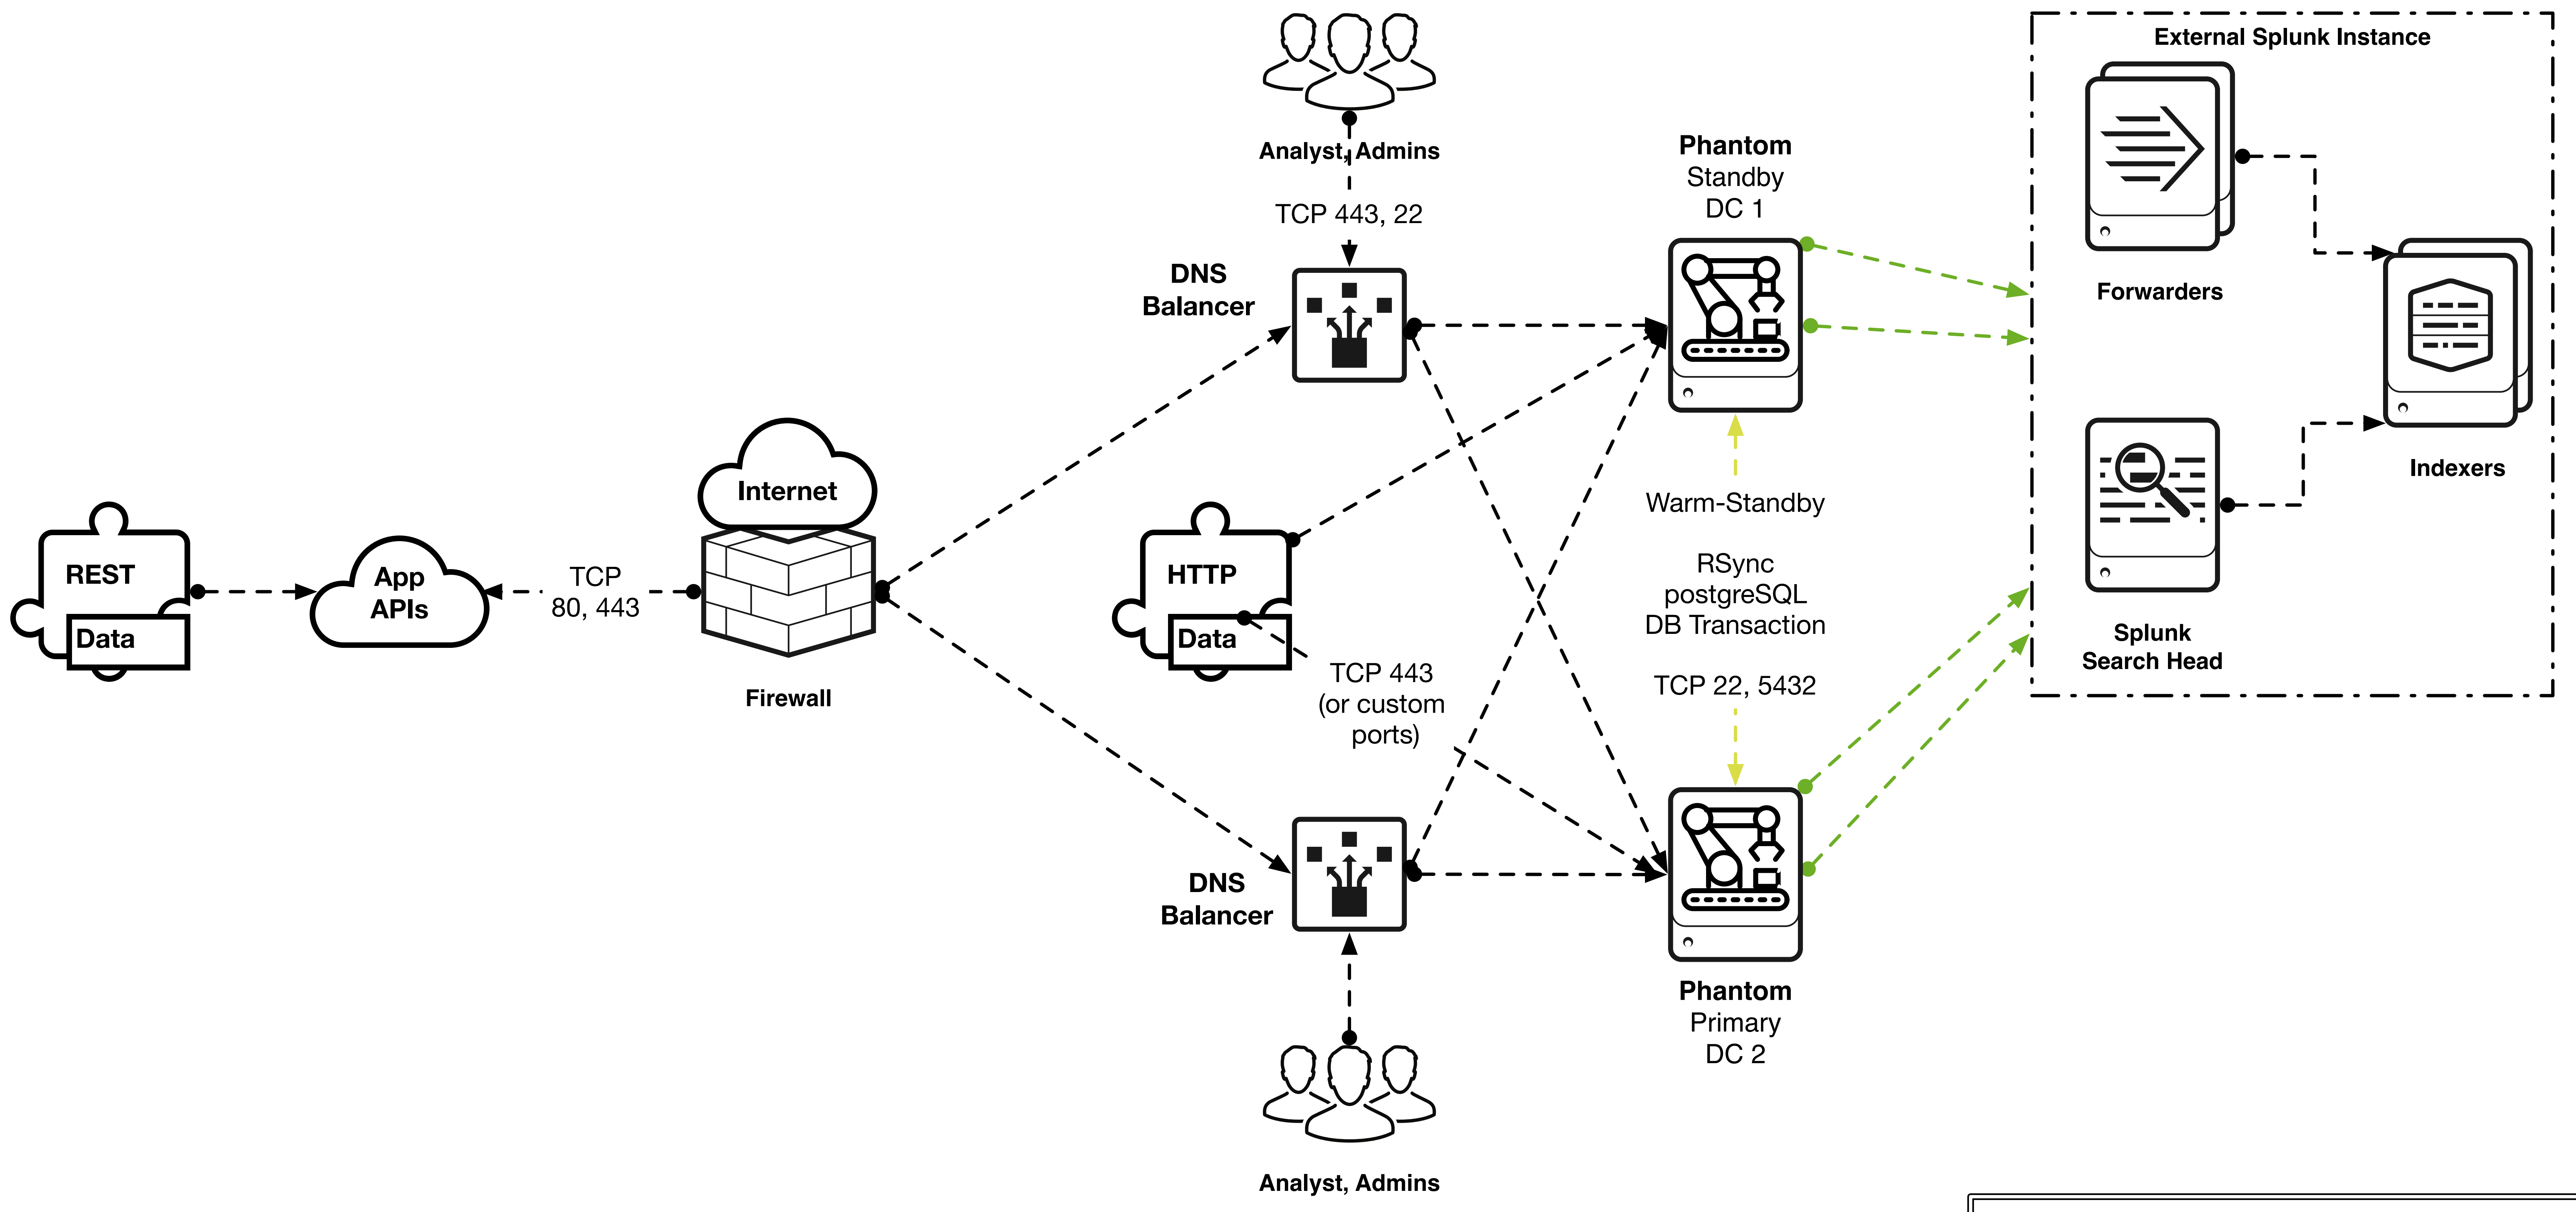

Warm Standby Process

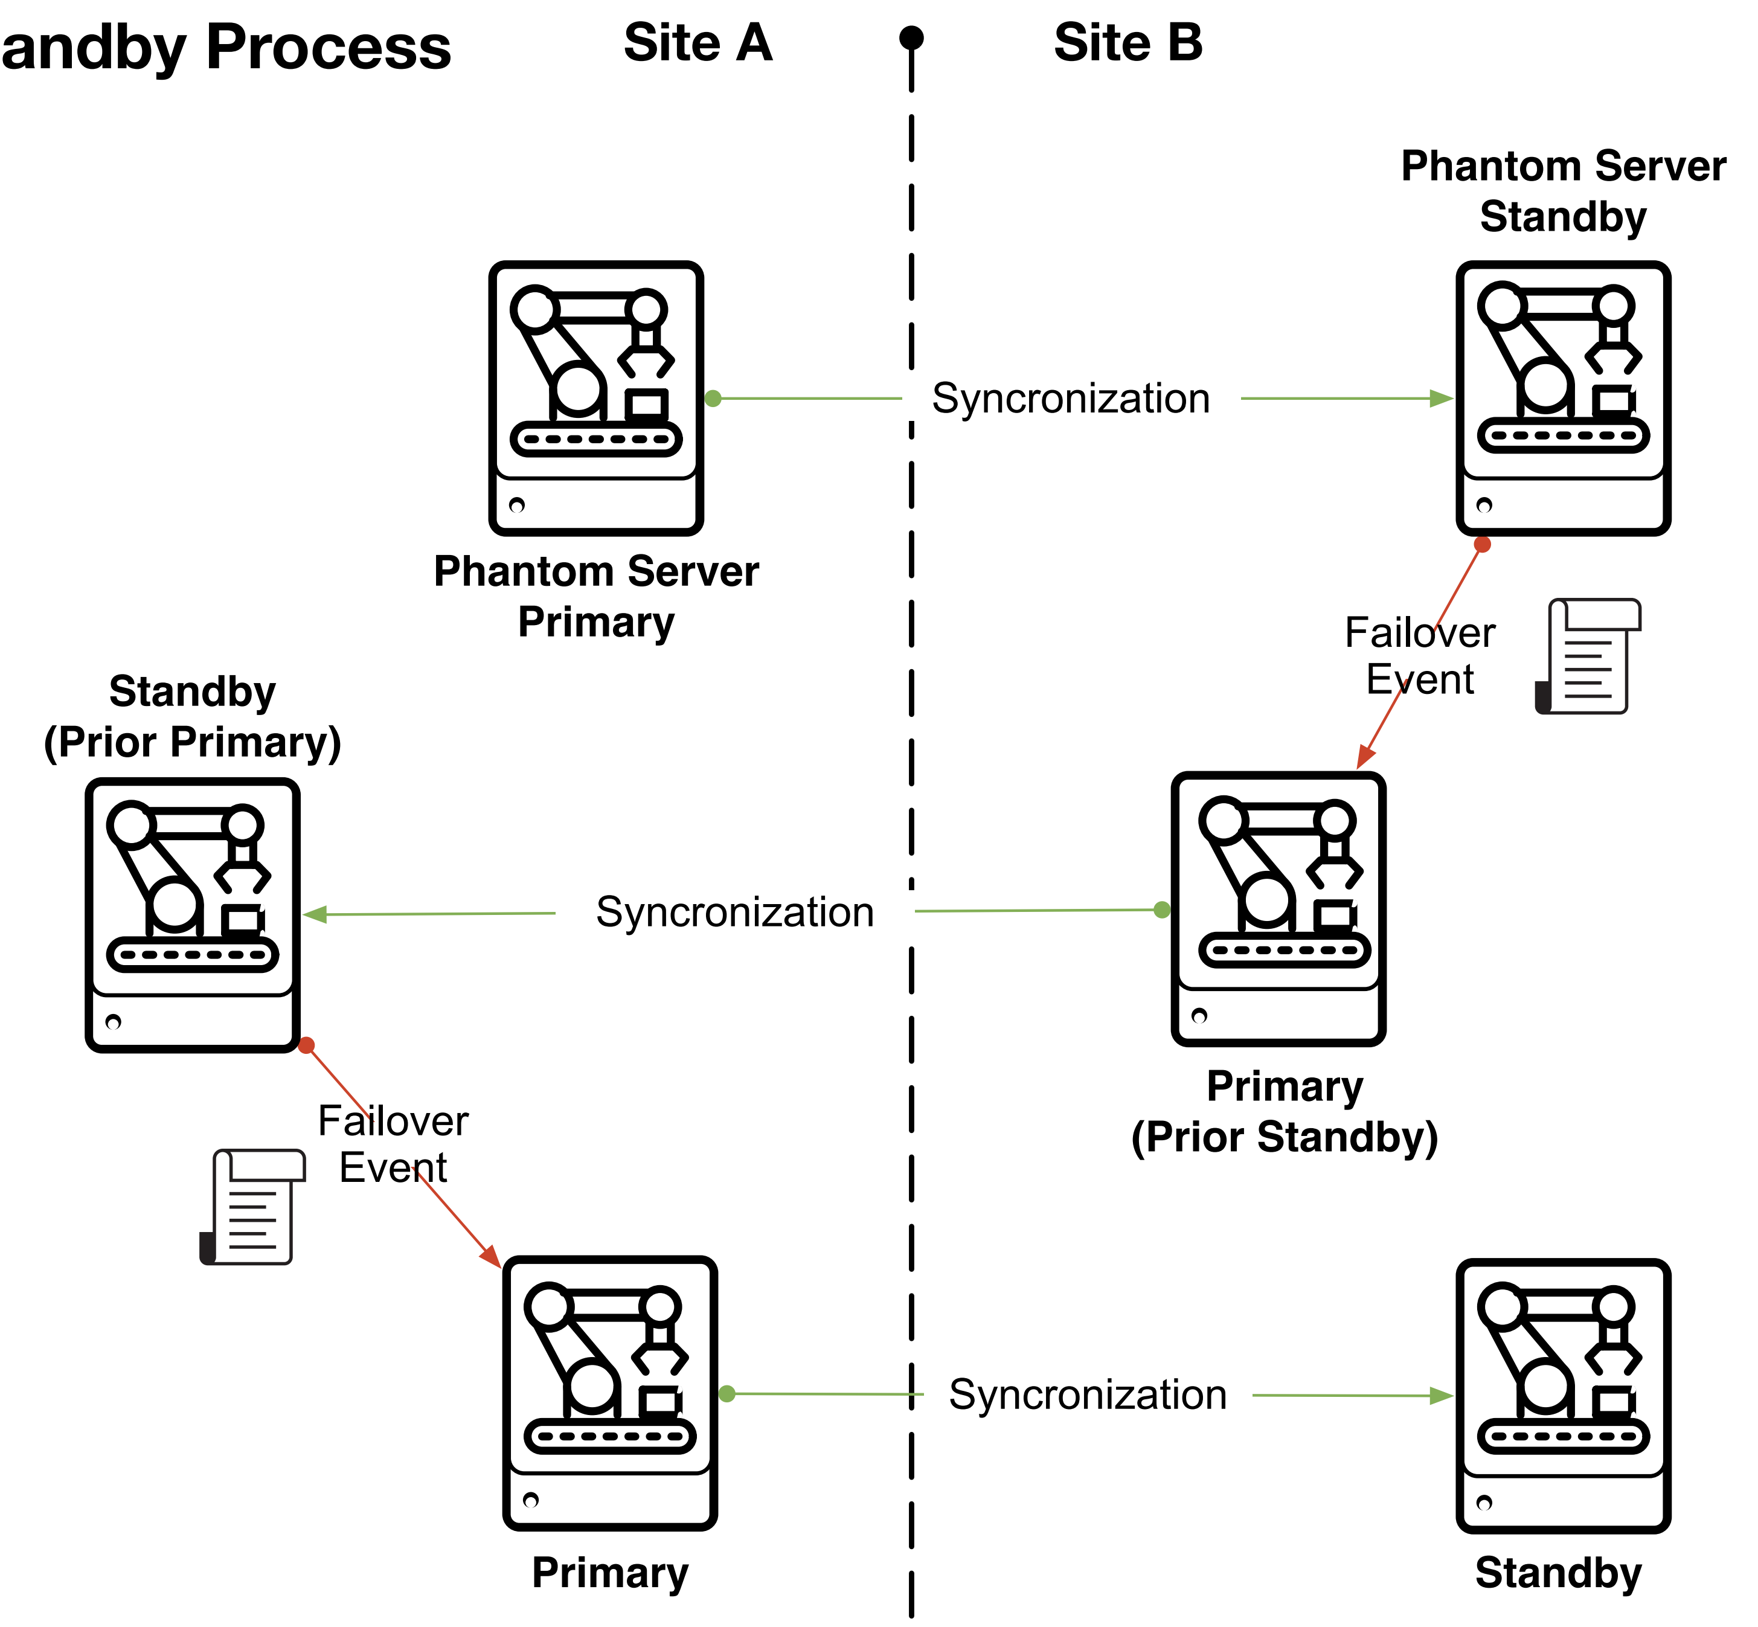

PROS:
Least number of resources
Native multi-site redundancy and high availability
Easiest to manage and gradual expansion increase process

CONS:
Performance constrained by number of available connections
Failover is scripted and can be semi-automated or manual

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation for Multiple Site Survivability

REVISIONS			
REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build capable for most loads < 1500 events per day and for teams using Case Management	3 MAR 19	

DRAWN BY	Architect
ISSUED TO	Customer

Company | **Company Confidential**

D2E+ - Distributed Warm Standby Deployment using embedded Splunk with customer provided Infrastructure

Intentionally left blank for diagram below

SpVA: D2E+

splunk> phantom

Splunk Phantom Stand-Alone Sizing Chart

Instance Type	Workloads with active playbooks	CPU cores	Memory GB
Large	>7000 events per hour	32	64
Medium	Up to 7000 events per hour	16	32
Small	Up to 4000 events per hour	8	16
Tiny	< 4000 events per hour	8	8

Recommended Sizing

Production Build Recommend Drive Mappings	Based on size (GB):	1024
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	204.8
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	245.8
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	307.2
/dev/mapper/centos-tmp	/tmp	51.2
/dev/mapper/centos-var	/var	51.2
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	5.1
/dev/mapper/centos-var_tmp	/var/tmp	51.2
/dev/mapper/centos-home	/home	102.4
	Total:	1018.9

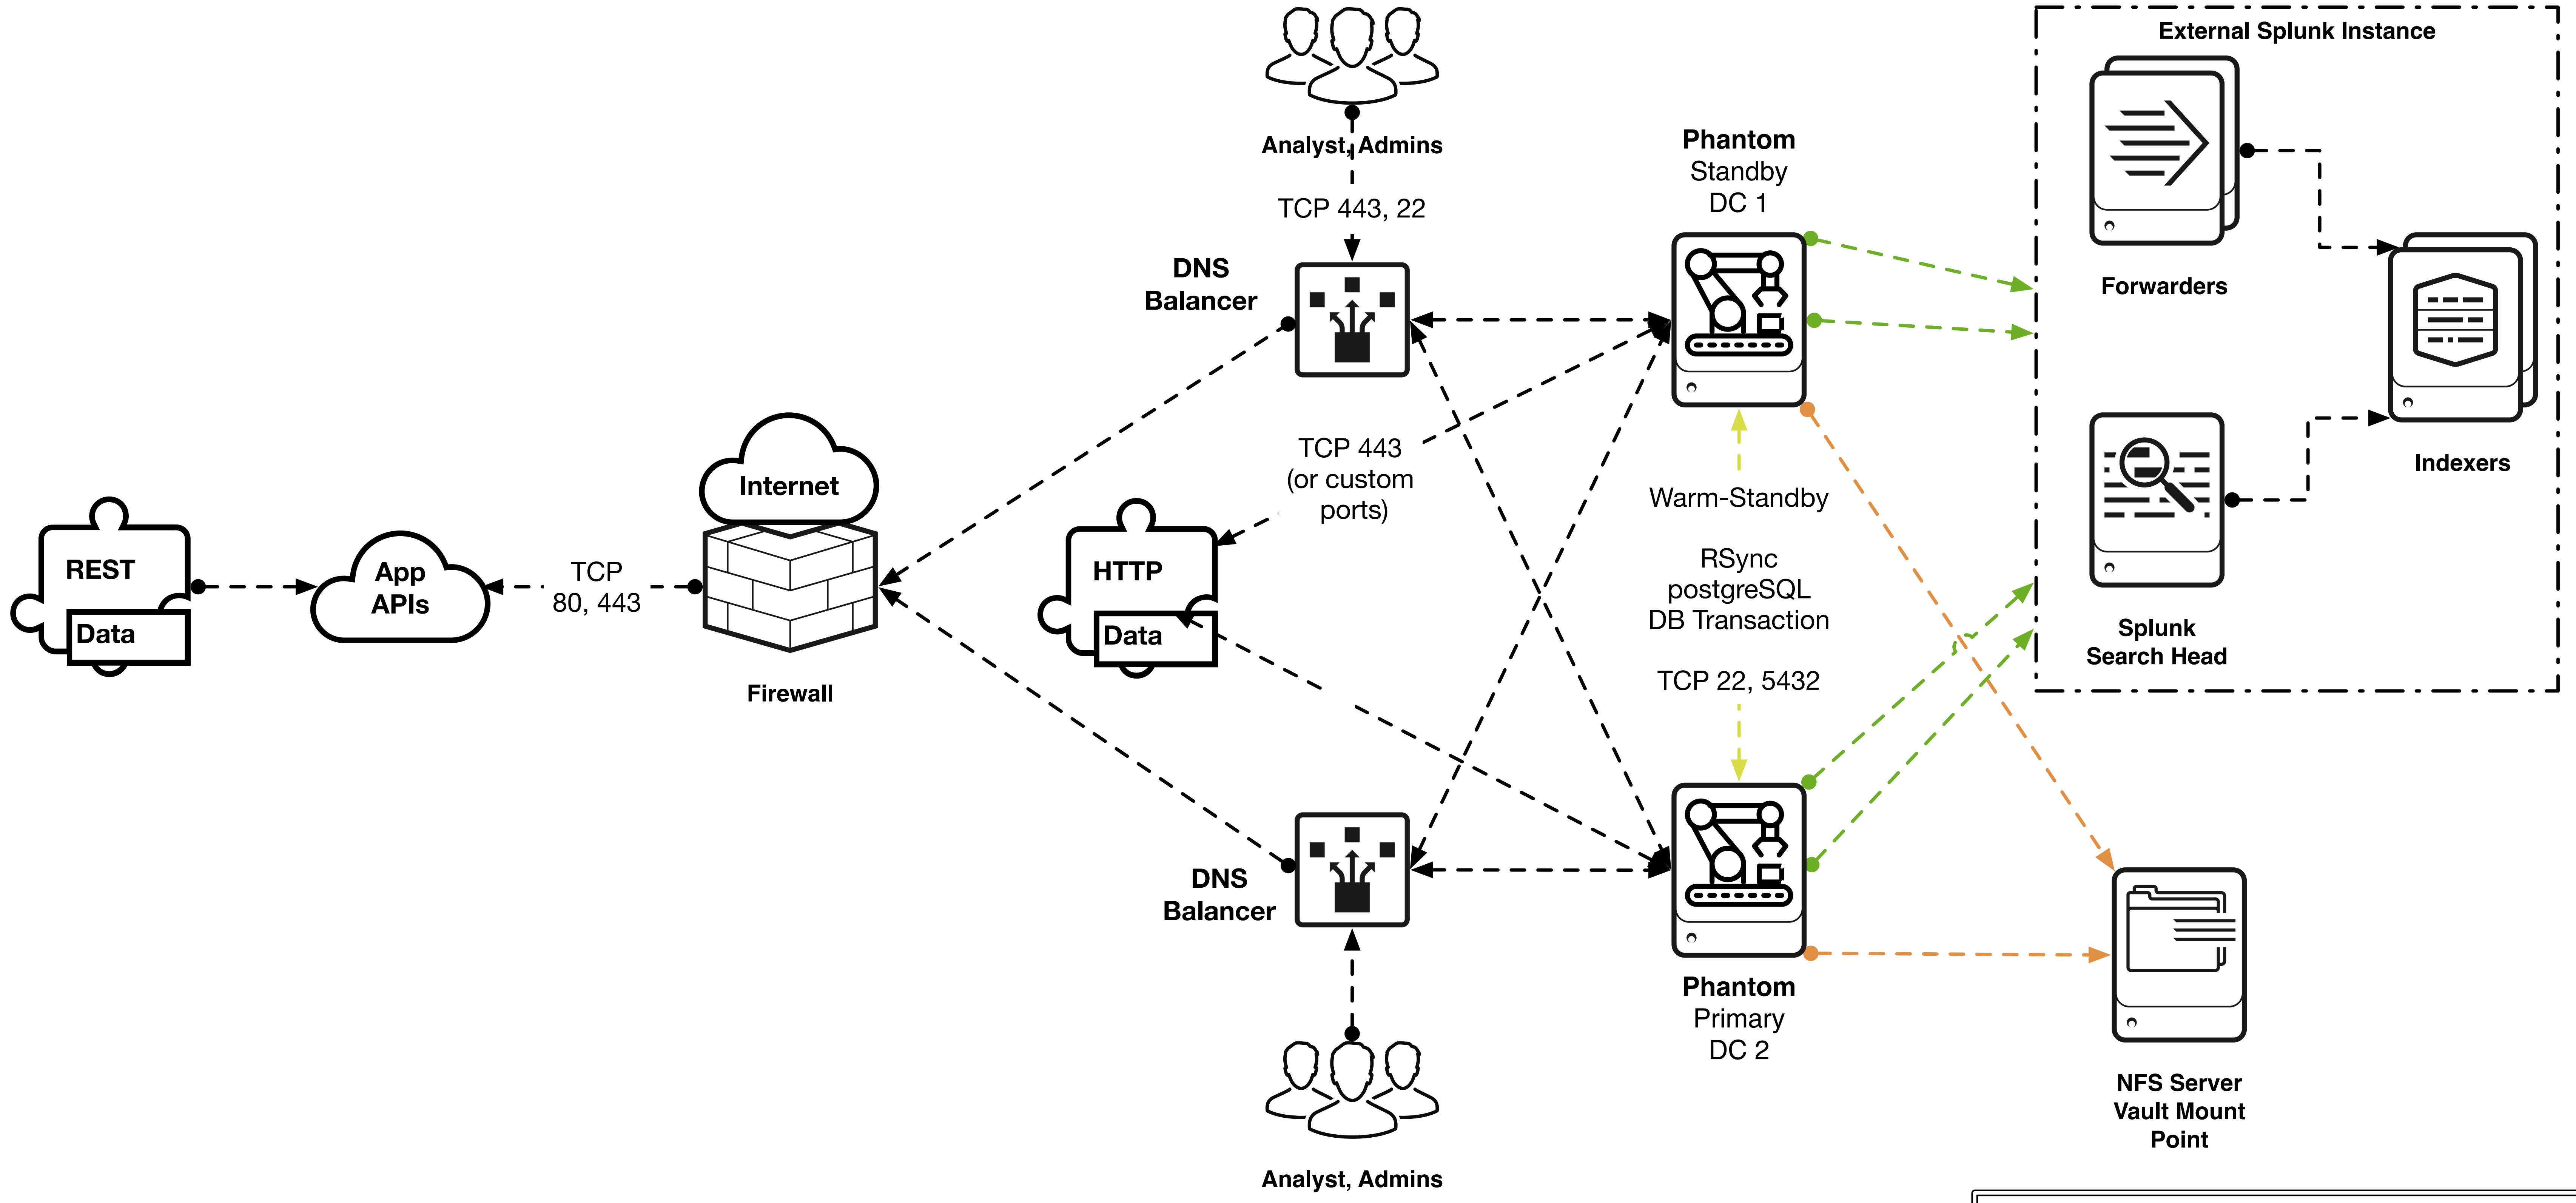

Warm Standby Process

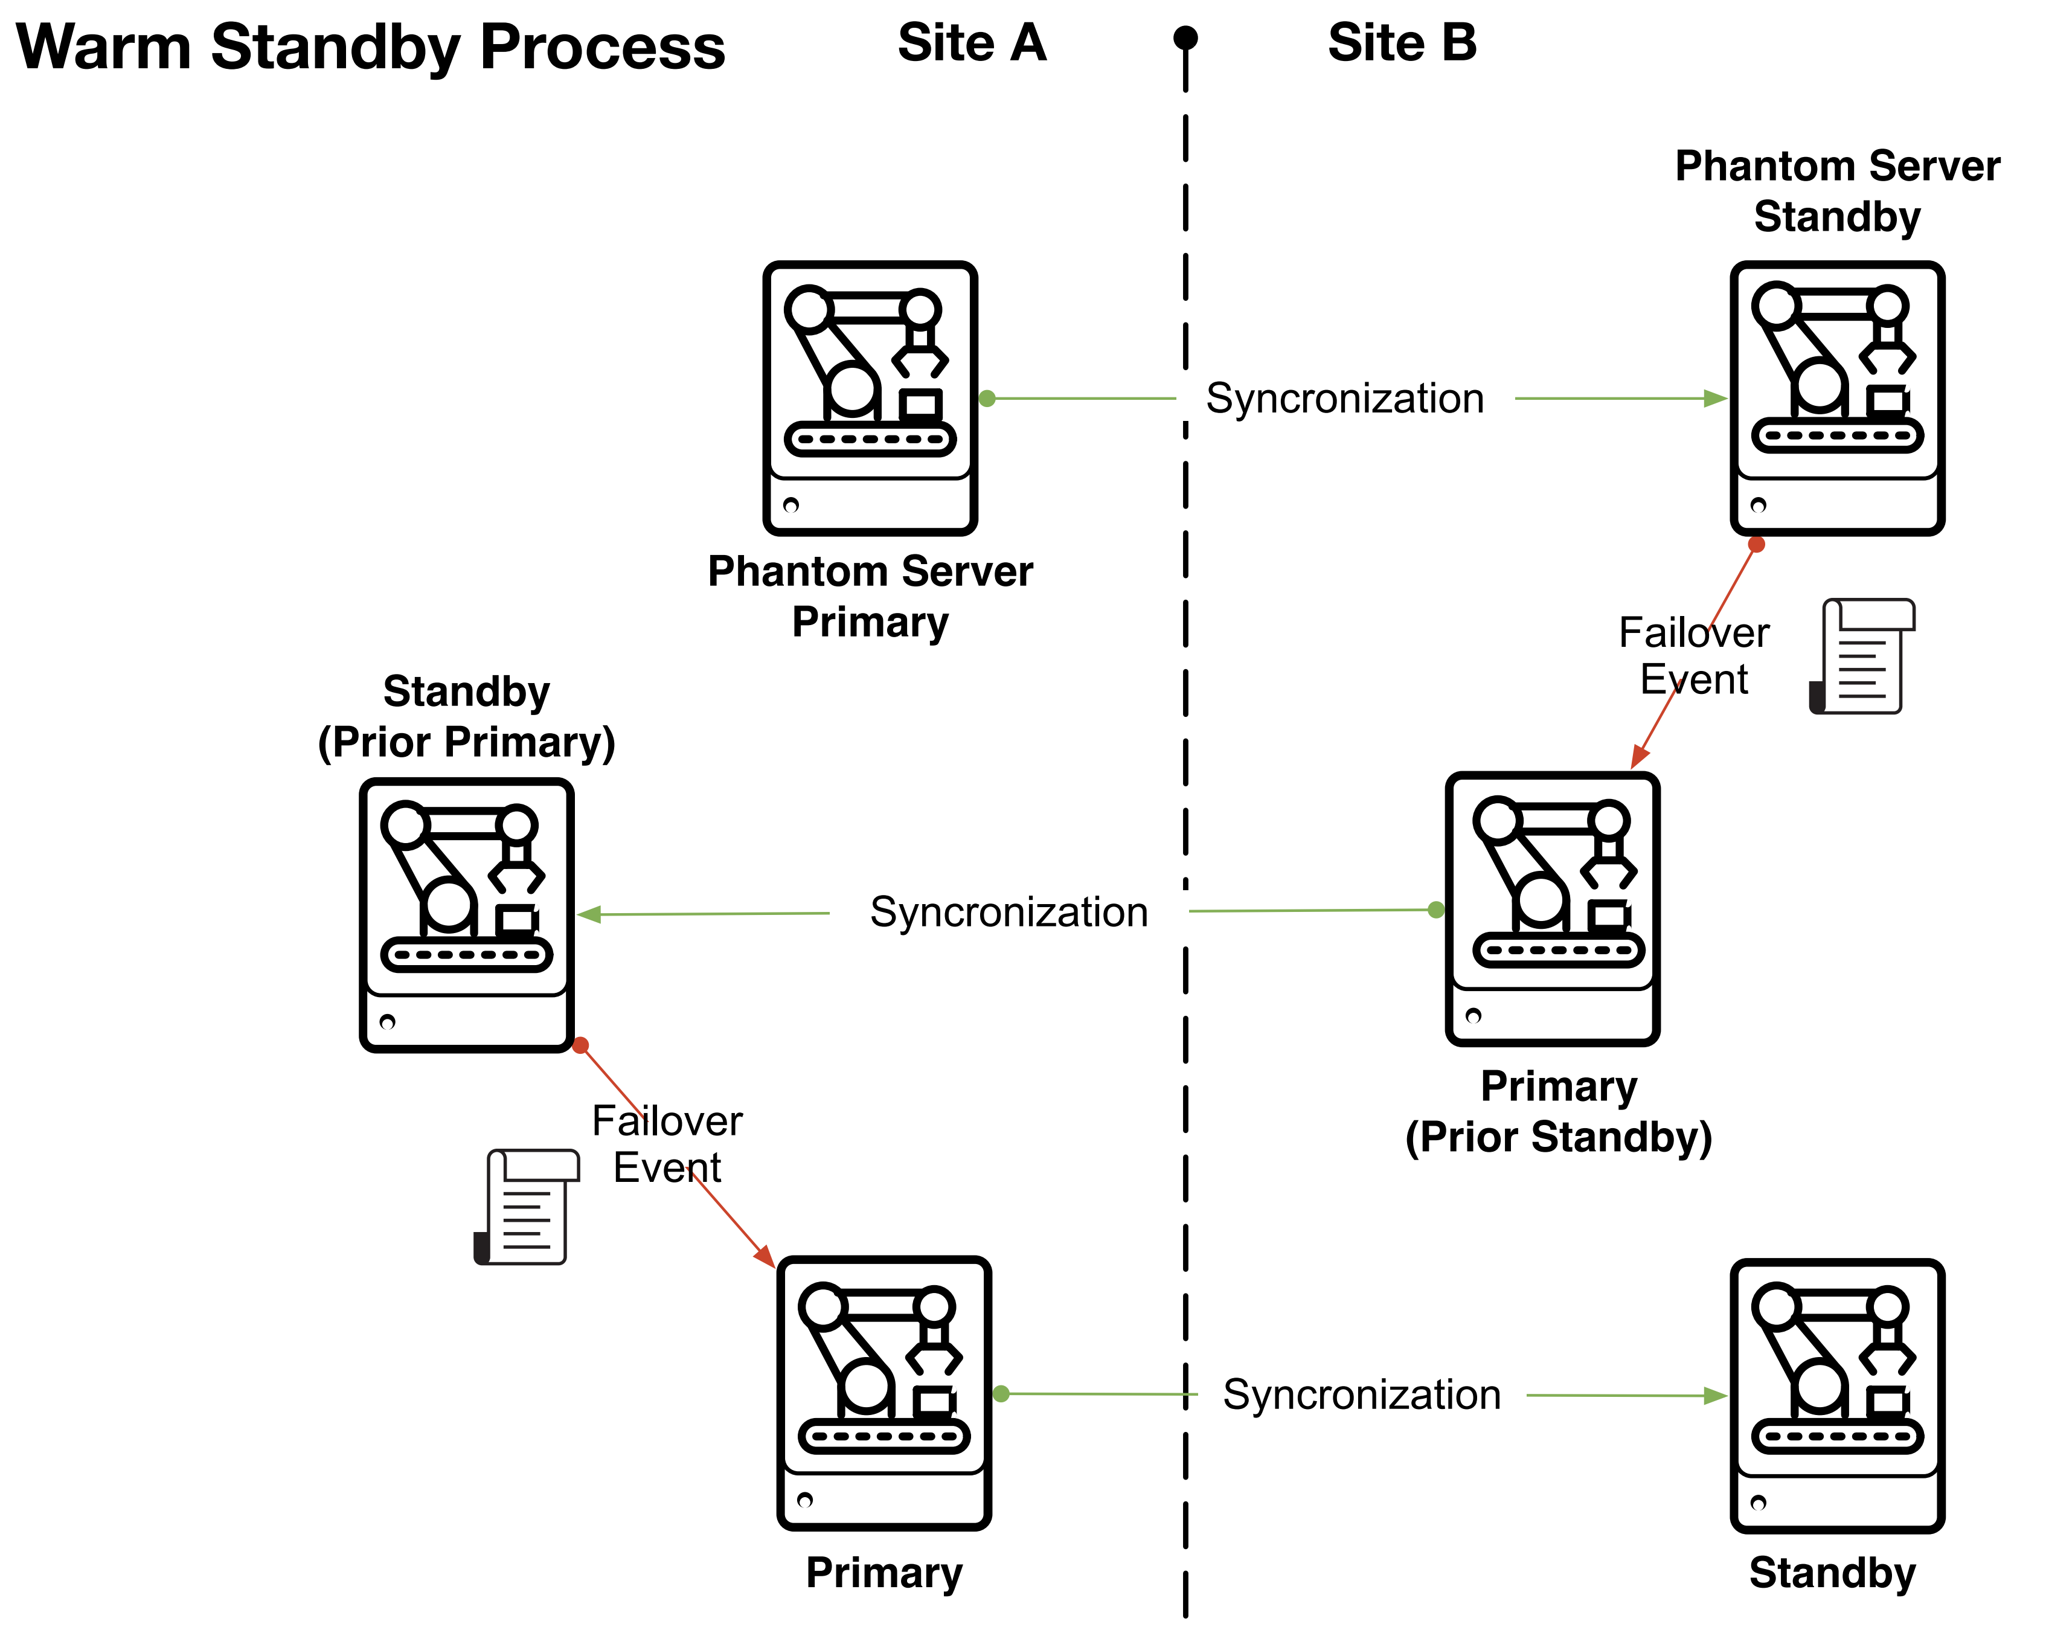

PROS:
Least number of resources
Native multi-site redundancy and high availability
Easiest to manage and gradual expansion increase process

CONS:
Performance constrained by number of available connections
Failover is scripted and can be semi-automated or manual

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation for Multiple Site Survivability

DRAWN BY	Architect
ISSUED TO	Customer

Company | **Company Confidential**

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build capable for most loads < 1500 events per day and for teams using Case Management	3 MAR 19	

D2CE - Distributed Warm Standby Deployment with Splunk Cloud Integration

Intentionally left blank for diagram below

SpVA: D2CE

splunk> phantom

Splunk Phantom Single Instance Sizing for Amazon Web Services

Instance Type	Workloads with an active playbook	EC2 Instance Sizing
Large	>7000 events per hour	c5.12xlarge
Medium Recommended Production from public website	Up to 7000 events per hour	c5.4xlarge
Small Recommended Development from public website Bare minimum from configuration	Up to 4000 events per hour	c5.2xlarge
Tiny ***Development ONLY***	< 4000 events per hour	c5.2xlarge

Recommended Sizing

Production Build Recommend Drive Mappings	Based on size (GB):	1024
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	204.8
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	245.8
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	307.2
/dev/mapper/centos-tmp	/tmp	51.2
/dev/mapper/centos-var	/var	51.2
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	5.1
/dev/mapper/centos-var_tmp	/var/tmp	51.2
/dev/mapper/centos-home	/home	102.4
	Total:	**1018.9**

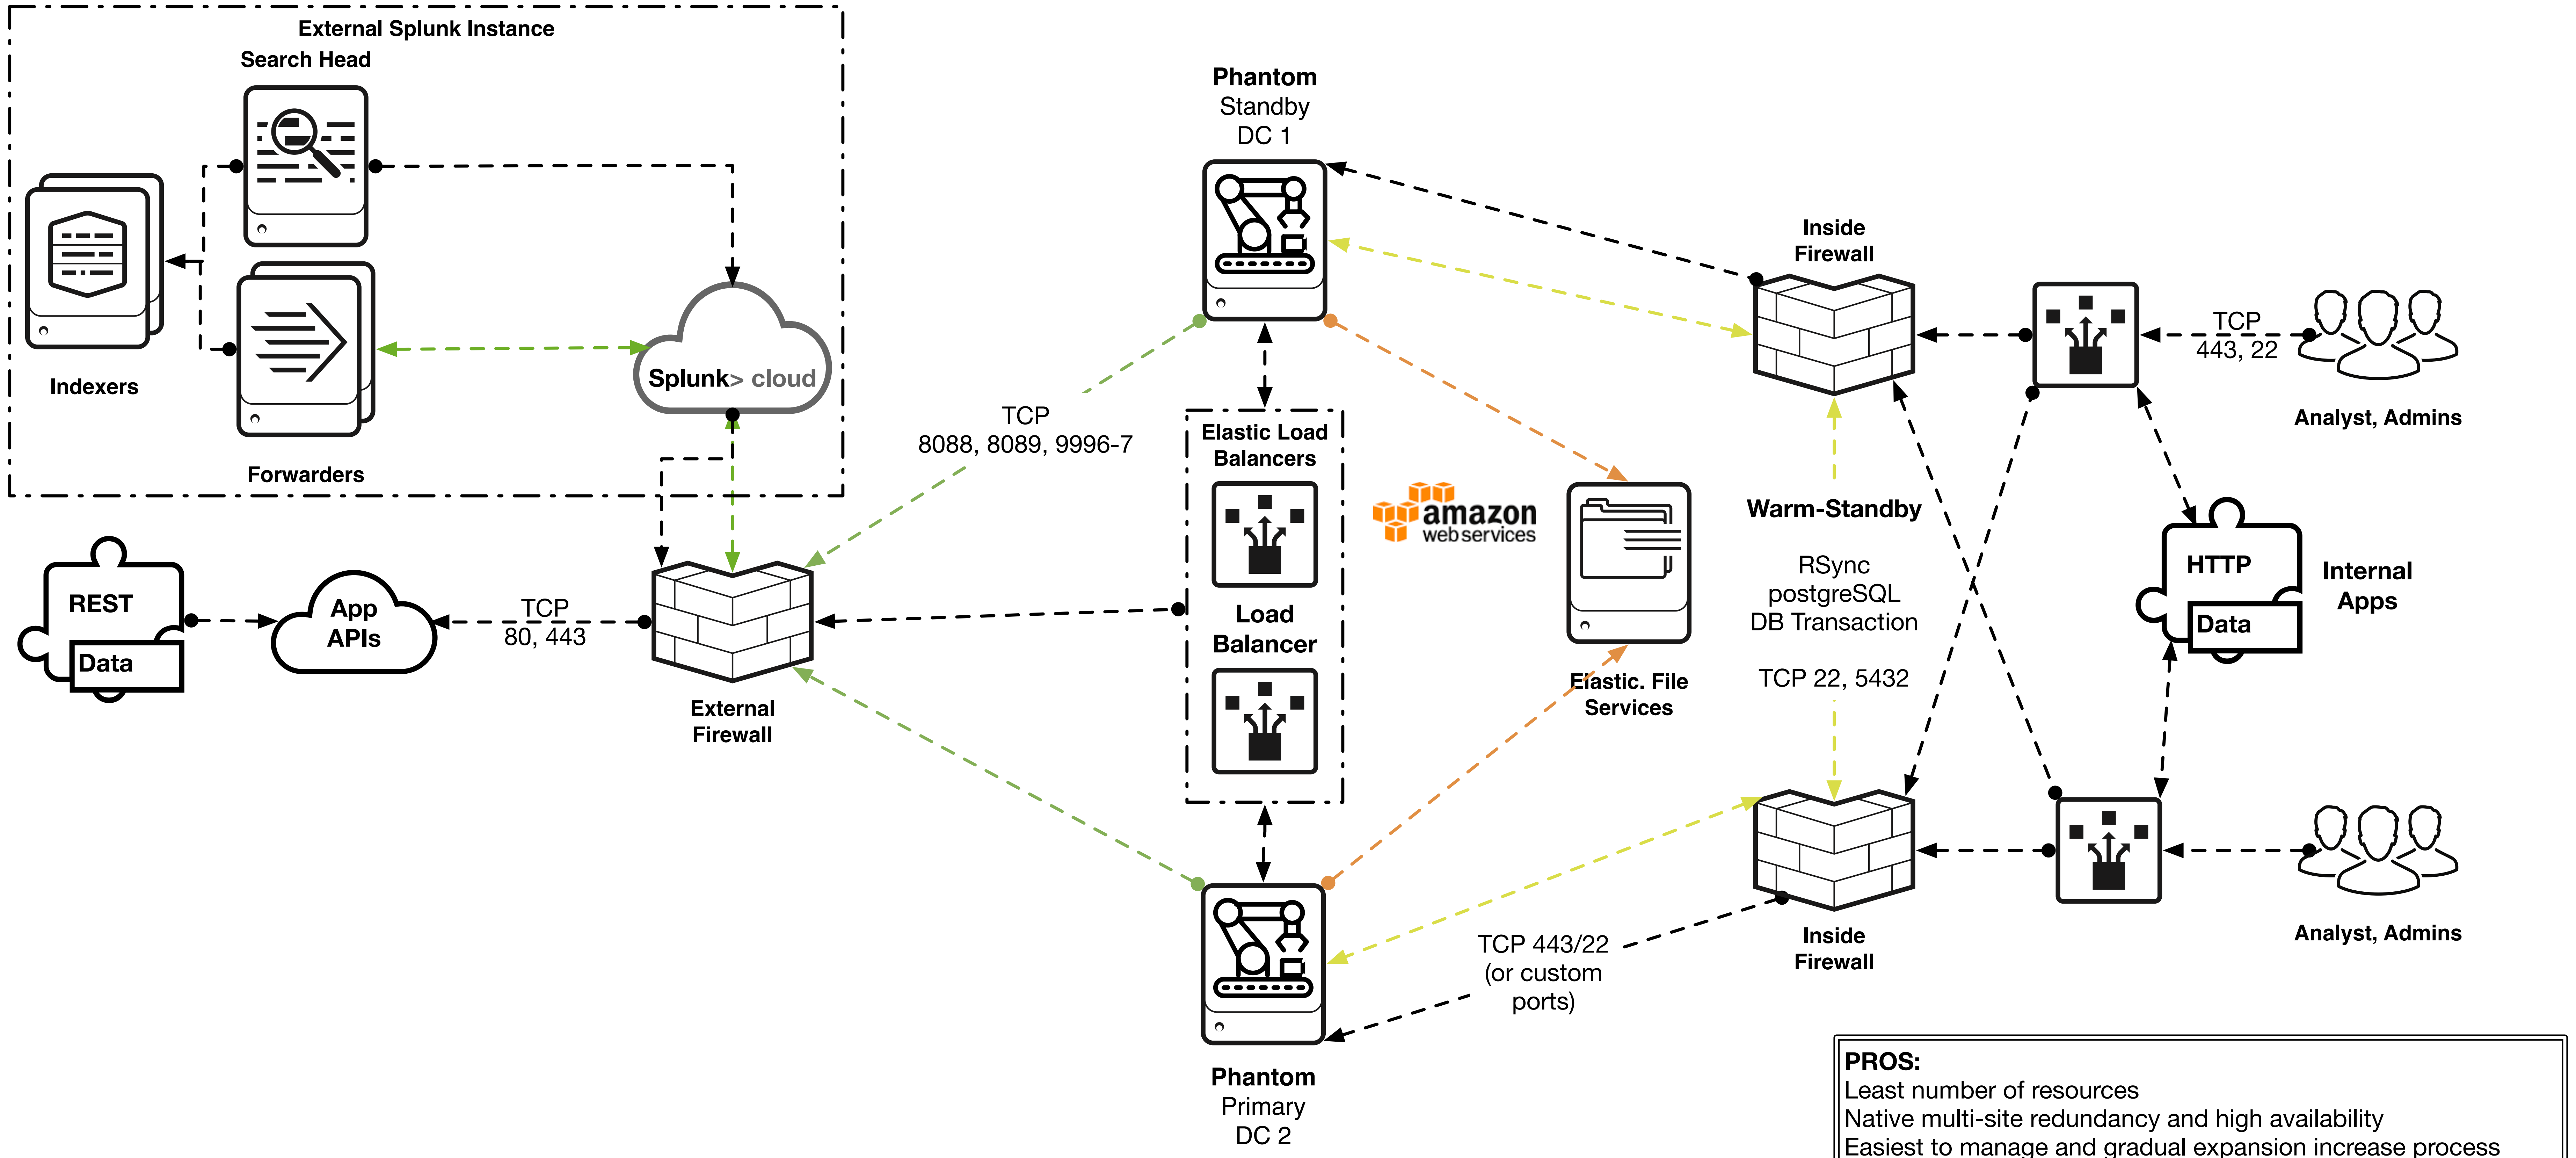

Warm Standby Process

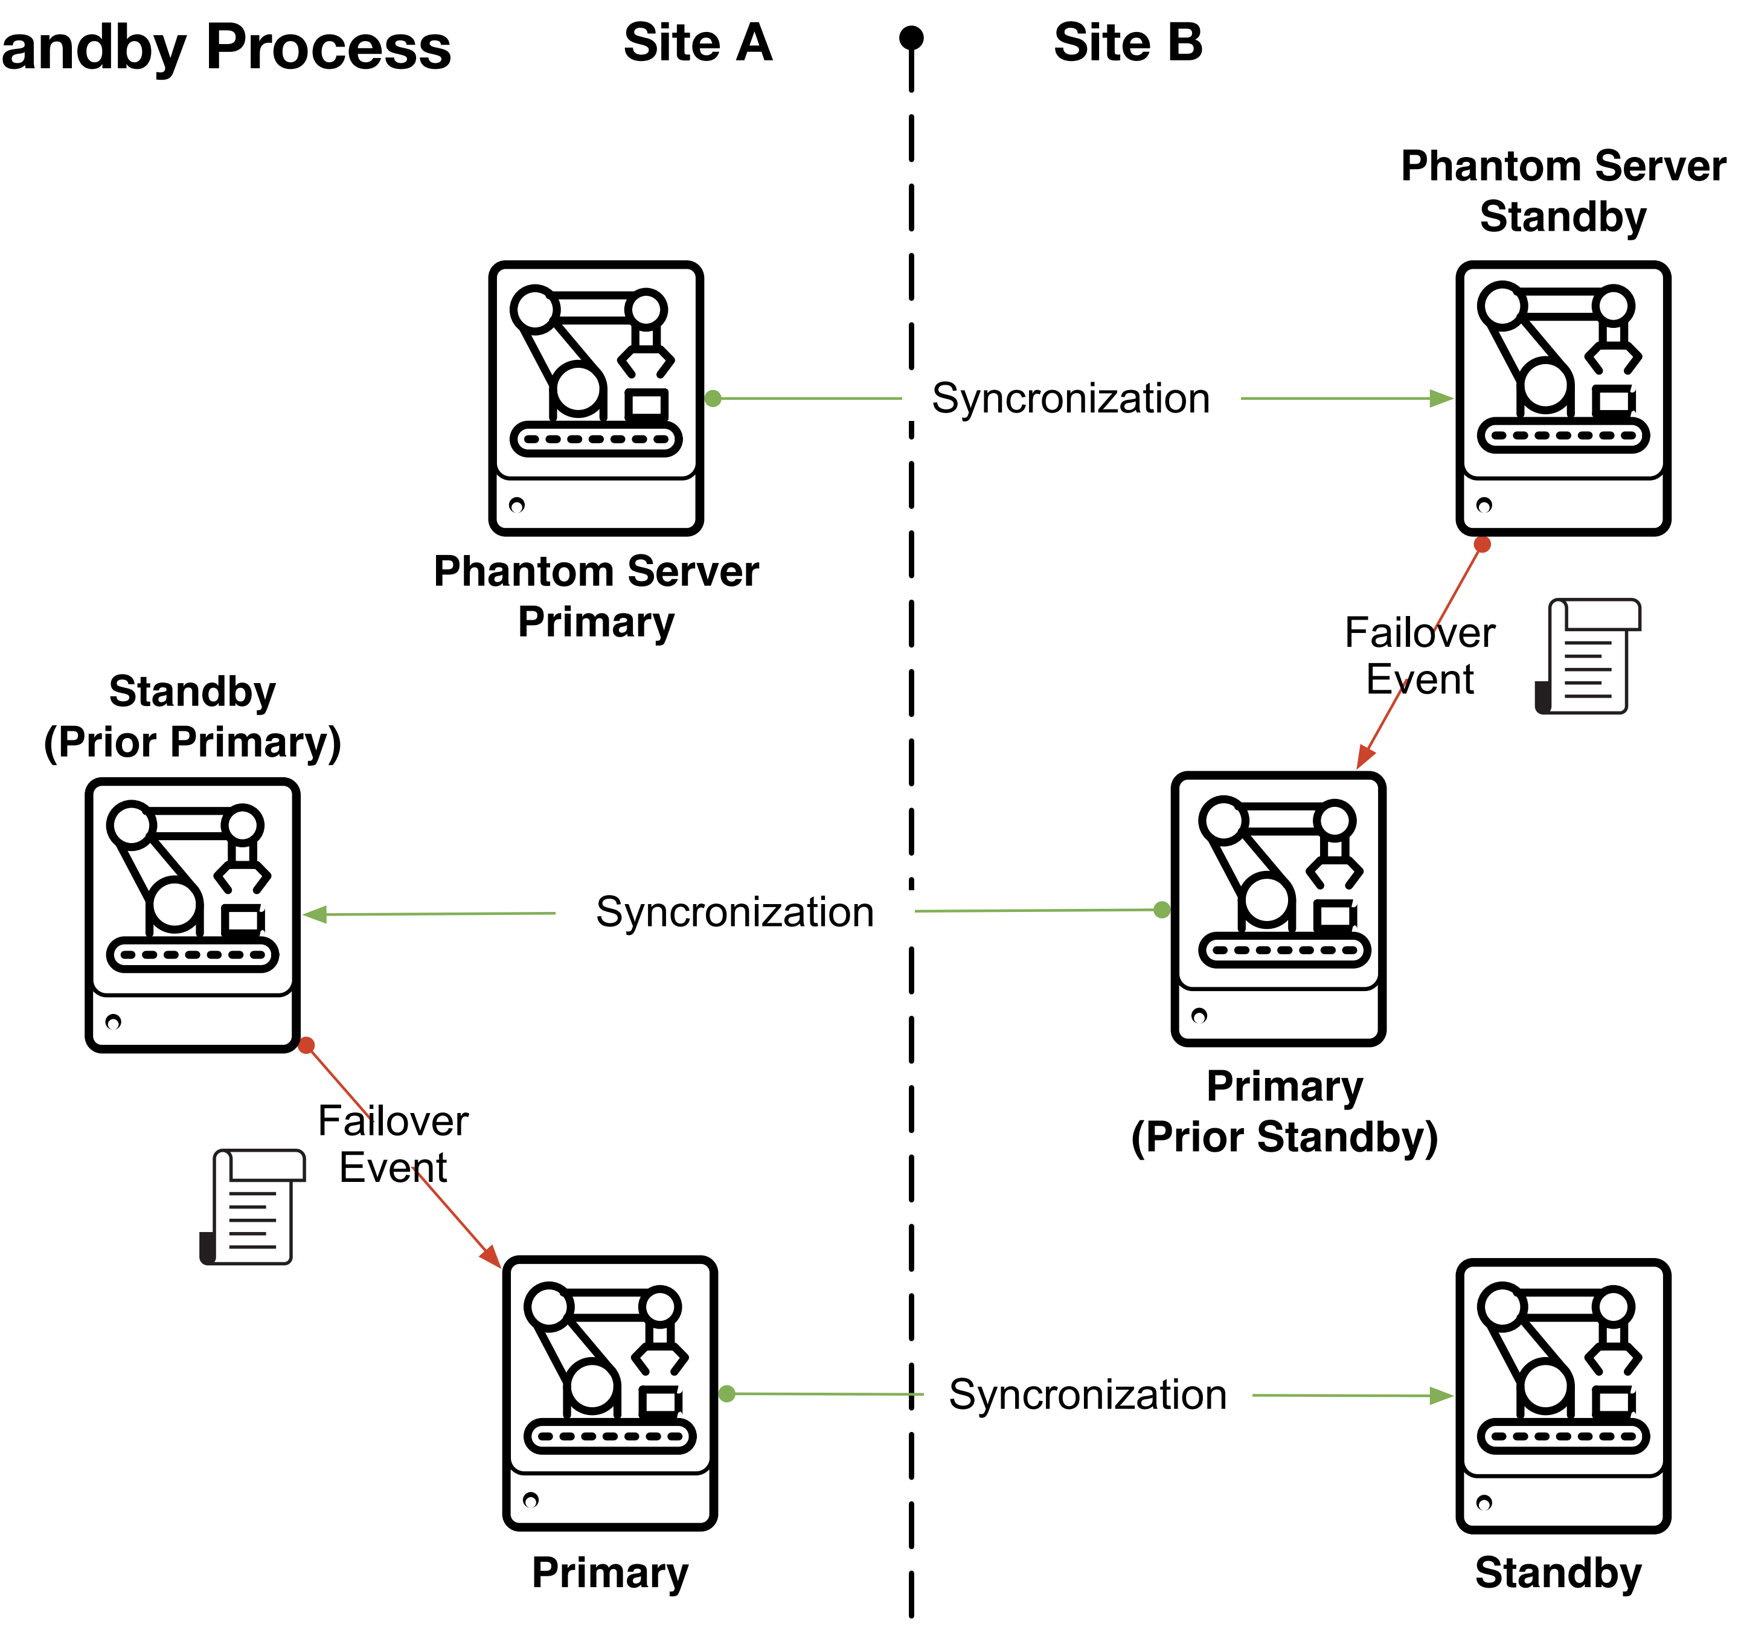

PROS:
Least number of resources
Native multi-site redundancy and high availability
Easiest to manage and gradual expansion increase process

CONS:
Performance constrained by number of available connections
Failover is scripted and can be semi-automated or manual

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 443	Used for sending Alerts to Phantom from Splunk> Cloud Phantom Add-On
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: **This is our Default PS Recommendation for Splunk Phantom and Cloud implementations. Phantom can be internalized (inside the DMZ), however some features will be reduced.**

DRAWN BY	Architect
ISSUED TO	Customer

Company | **Company Confidential**

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build capable for most loads < 1500 events per day and for teams using Case Management	3 MAR 19	
1.2	Default access to Phantom on Internal infrastructure. Splunk Cloud must have the Phantom app and API access. **Phantom is in the DMZ or DMZ like**	26 NOV 19	
1.3	Updated ports and localized External Splunk architecture	27 APR 20	

D2CE+ - Distributed Warm Standby Deployment with Splunk Cloud Integration

Intentionally left blank for diagram below

SpVA: D2CE+

splunk> phantom

Splunk Phantom Single Instance Sizing for Amazon Web Services

Instance Type	Workloads with an active playbook	EC2 Instance Sizing
Large	>7000 events per hour	c5.12xlarge
Medium Recommended Production from public website	Up to 7000 events per hour	c5.4xlarge
Small Recommended Development from public website Bare minimum from configuration	Up to 4000 events per hour	c5.2xlarge
Tiny *Development ONLY*	< 4000 events per hour	c5.2xlarge

Recommended Sizing

Production Build Recommend Drive Mappings	Based on size (GB):	1024
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	204.8
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	245.8
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	307.2
/dev/mapper/centos-tmp	/tmp	51.2
/dev/mapper/centos-var	/var	51.2
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	5.1
/dev/mapper/centos-var_tmp	/var/tmp	51.2
/dev/mapper/centos-home	/home	102.4
	Total:	1018.9

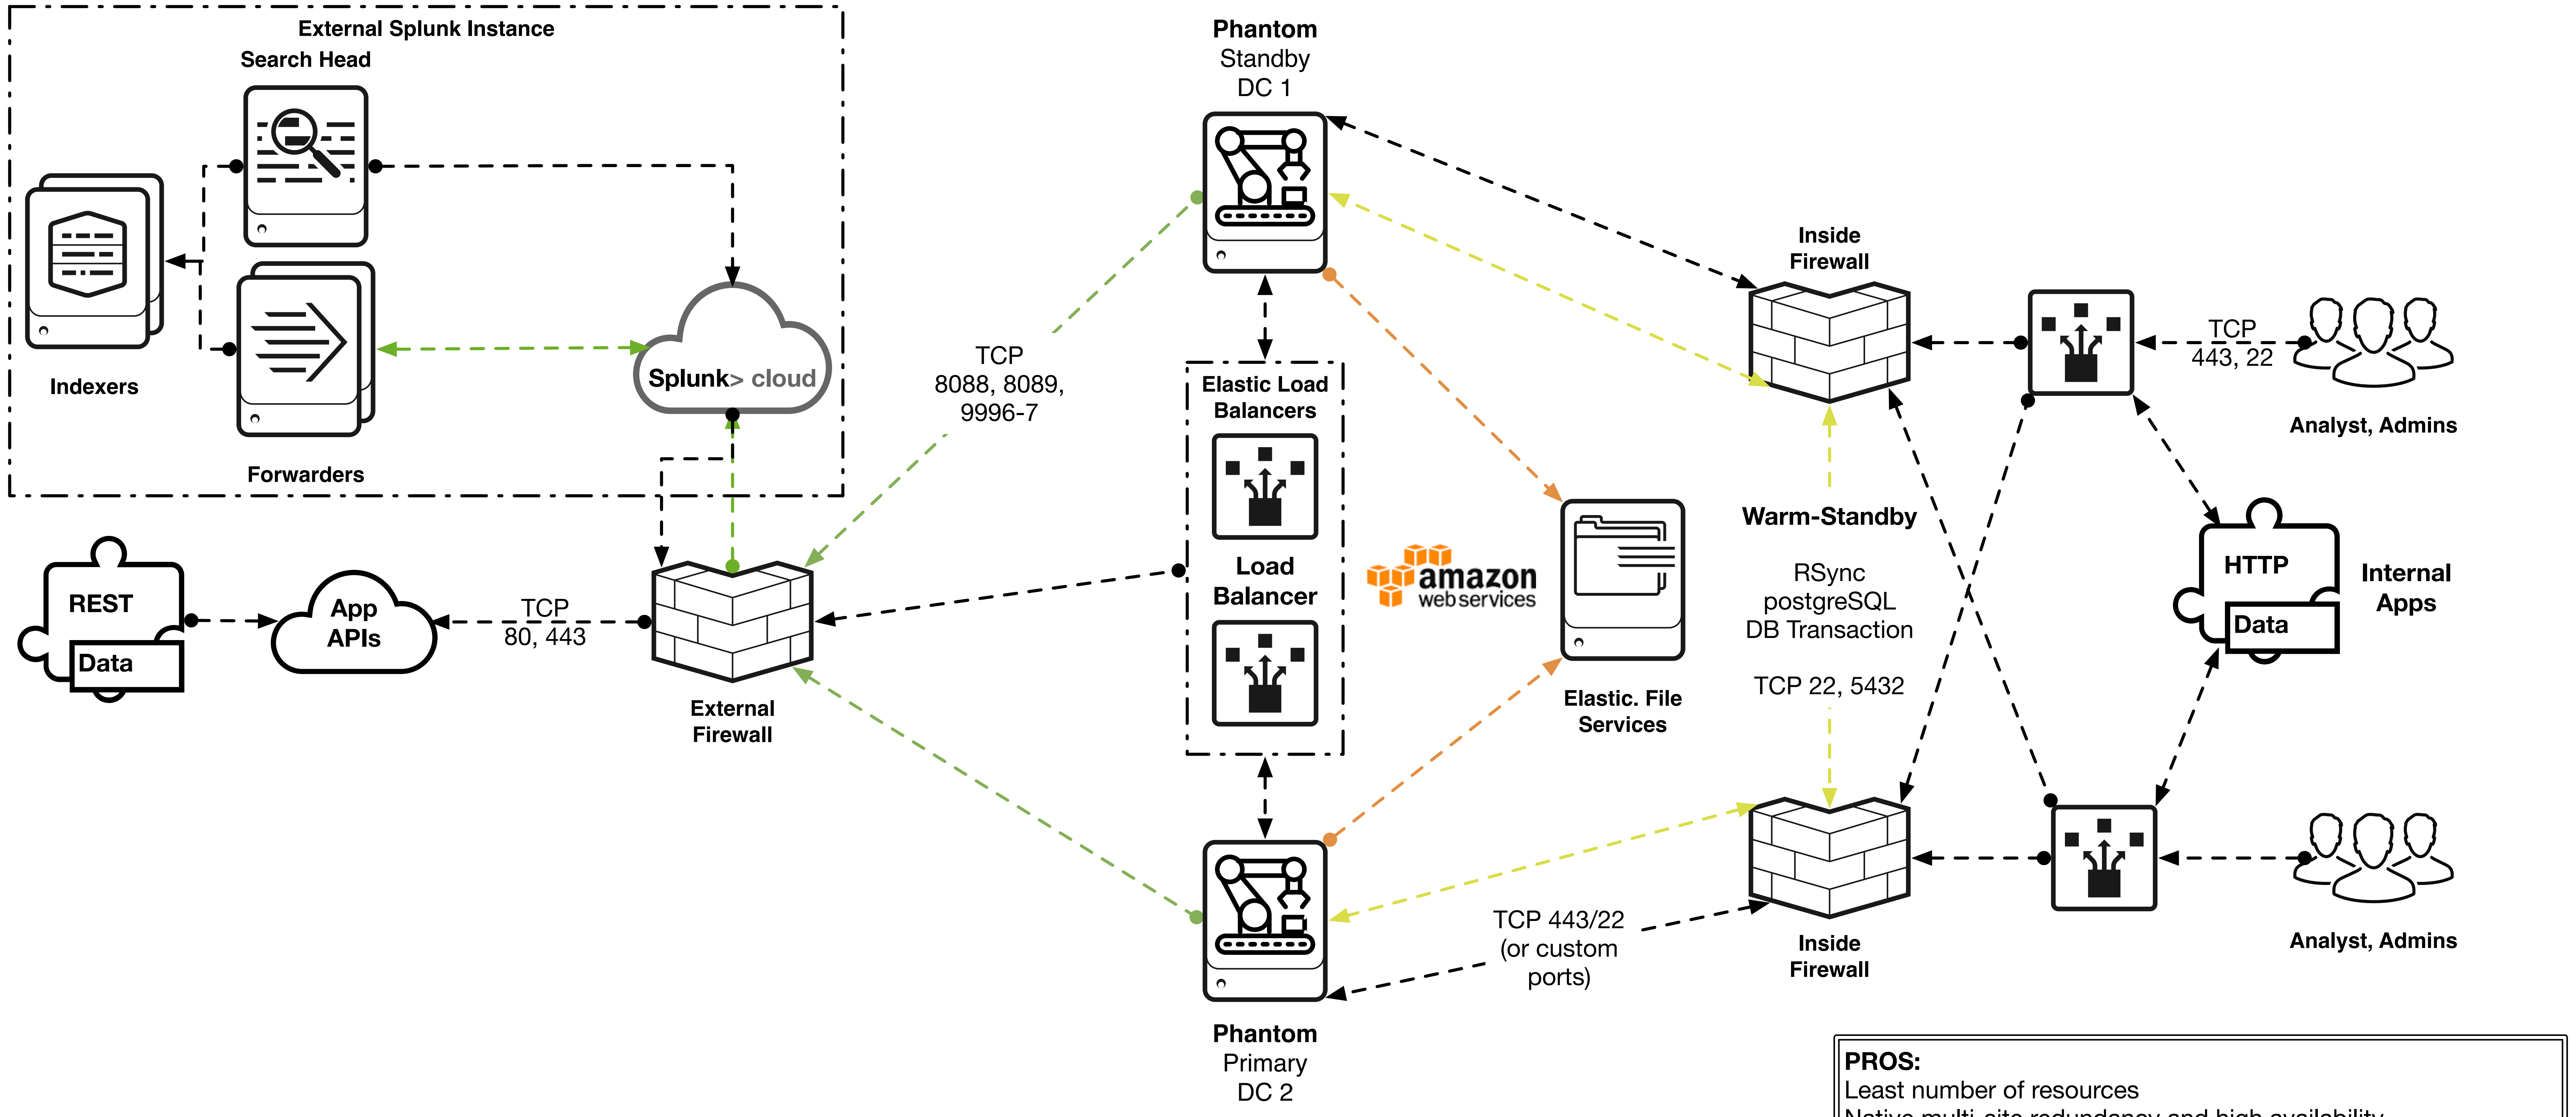

Warm Standby Process

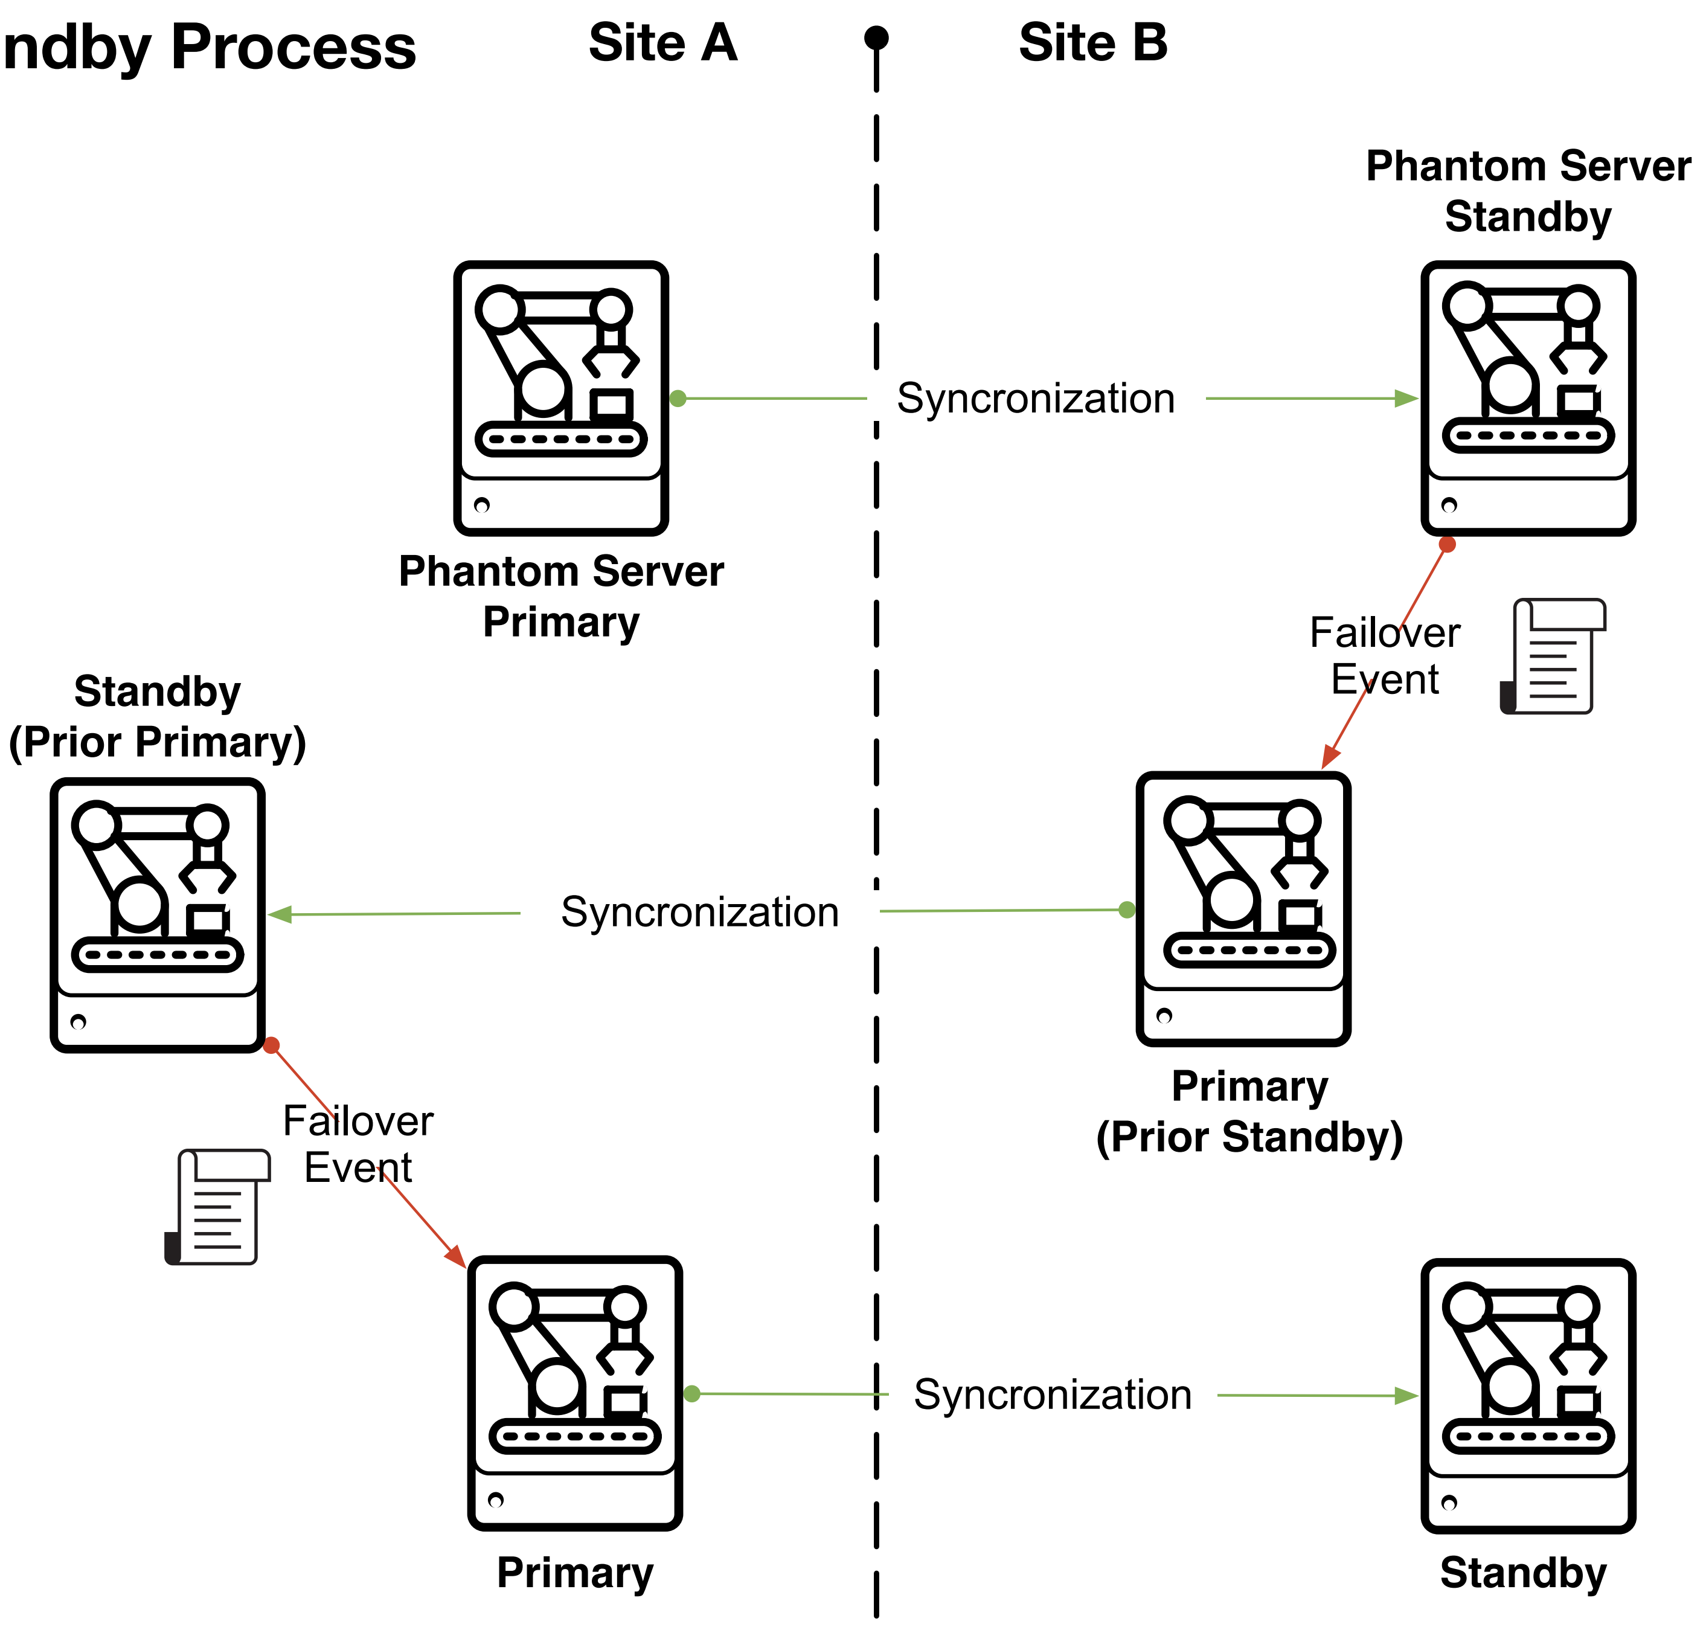

PROS:
Least number of resources
Native multi-site redundancy and high availability
Easiest to manage and gradual expansion increase process

CONS:
Performance constrained by number of available connections
Failover is scripted and can be semi-automated or manual

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 443	Used for sending Alerts to Phantom from Splunk> Cloud Phantom Add-On
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

Standalone Instance

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation for Splunk Phantom and Cloud implementations. Phantom can be internalized (inside the DMZ), however some features will be reduced.

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build capable for most loads < 1500 events per day and for teams using Case Management	3 MAR 19	
1.2	Default access to Phantom on Internal infrastructure. Splunk Cloud must have the Phantom app and API access. **Phantom is in the DMZ or DMZ like**	26 NOV 19	
1.3	Updated ports and localized External Splunk architecture	27 APR 20	

DRAWN BY	Architect
ISSUED TO	Customer

Company | **Company Confidential**

C1 - High Capacity Clustered Deployment

Intentionally left blank for diagram below

Splunk Phantom Cluster Sizing Plan

Cluster Type	Workloads with active playbooks	DB/ Common Node CPU cores	DB/Common Node Memory GB	Phantom Node CPU cores	Phantom Node Memory GB	Number of Phantom Nodes
XLarge	>50,000	32	64	16-32	32	8
Large	Up to 25,000-50,000 per hour	16	64	8	16	8
Medium	Up to 25,000 events per hour	16	32	8	16	5
Small	Up to 10,000 events per hour	8	32	4	8	3

Recommended Sizing

SpVA: C1

splunk> phantom

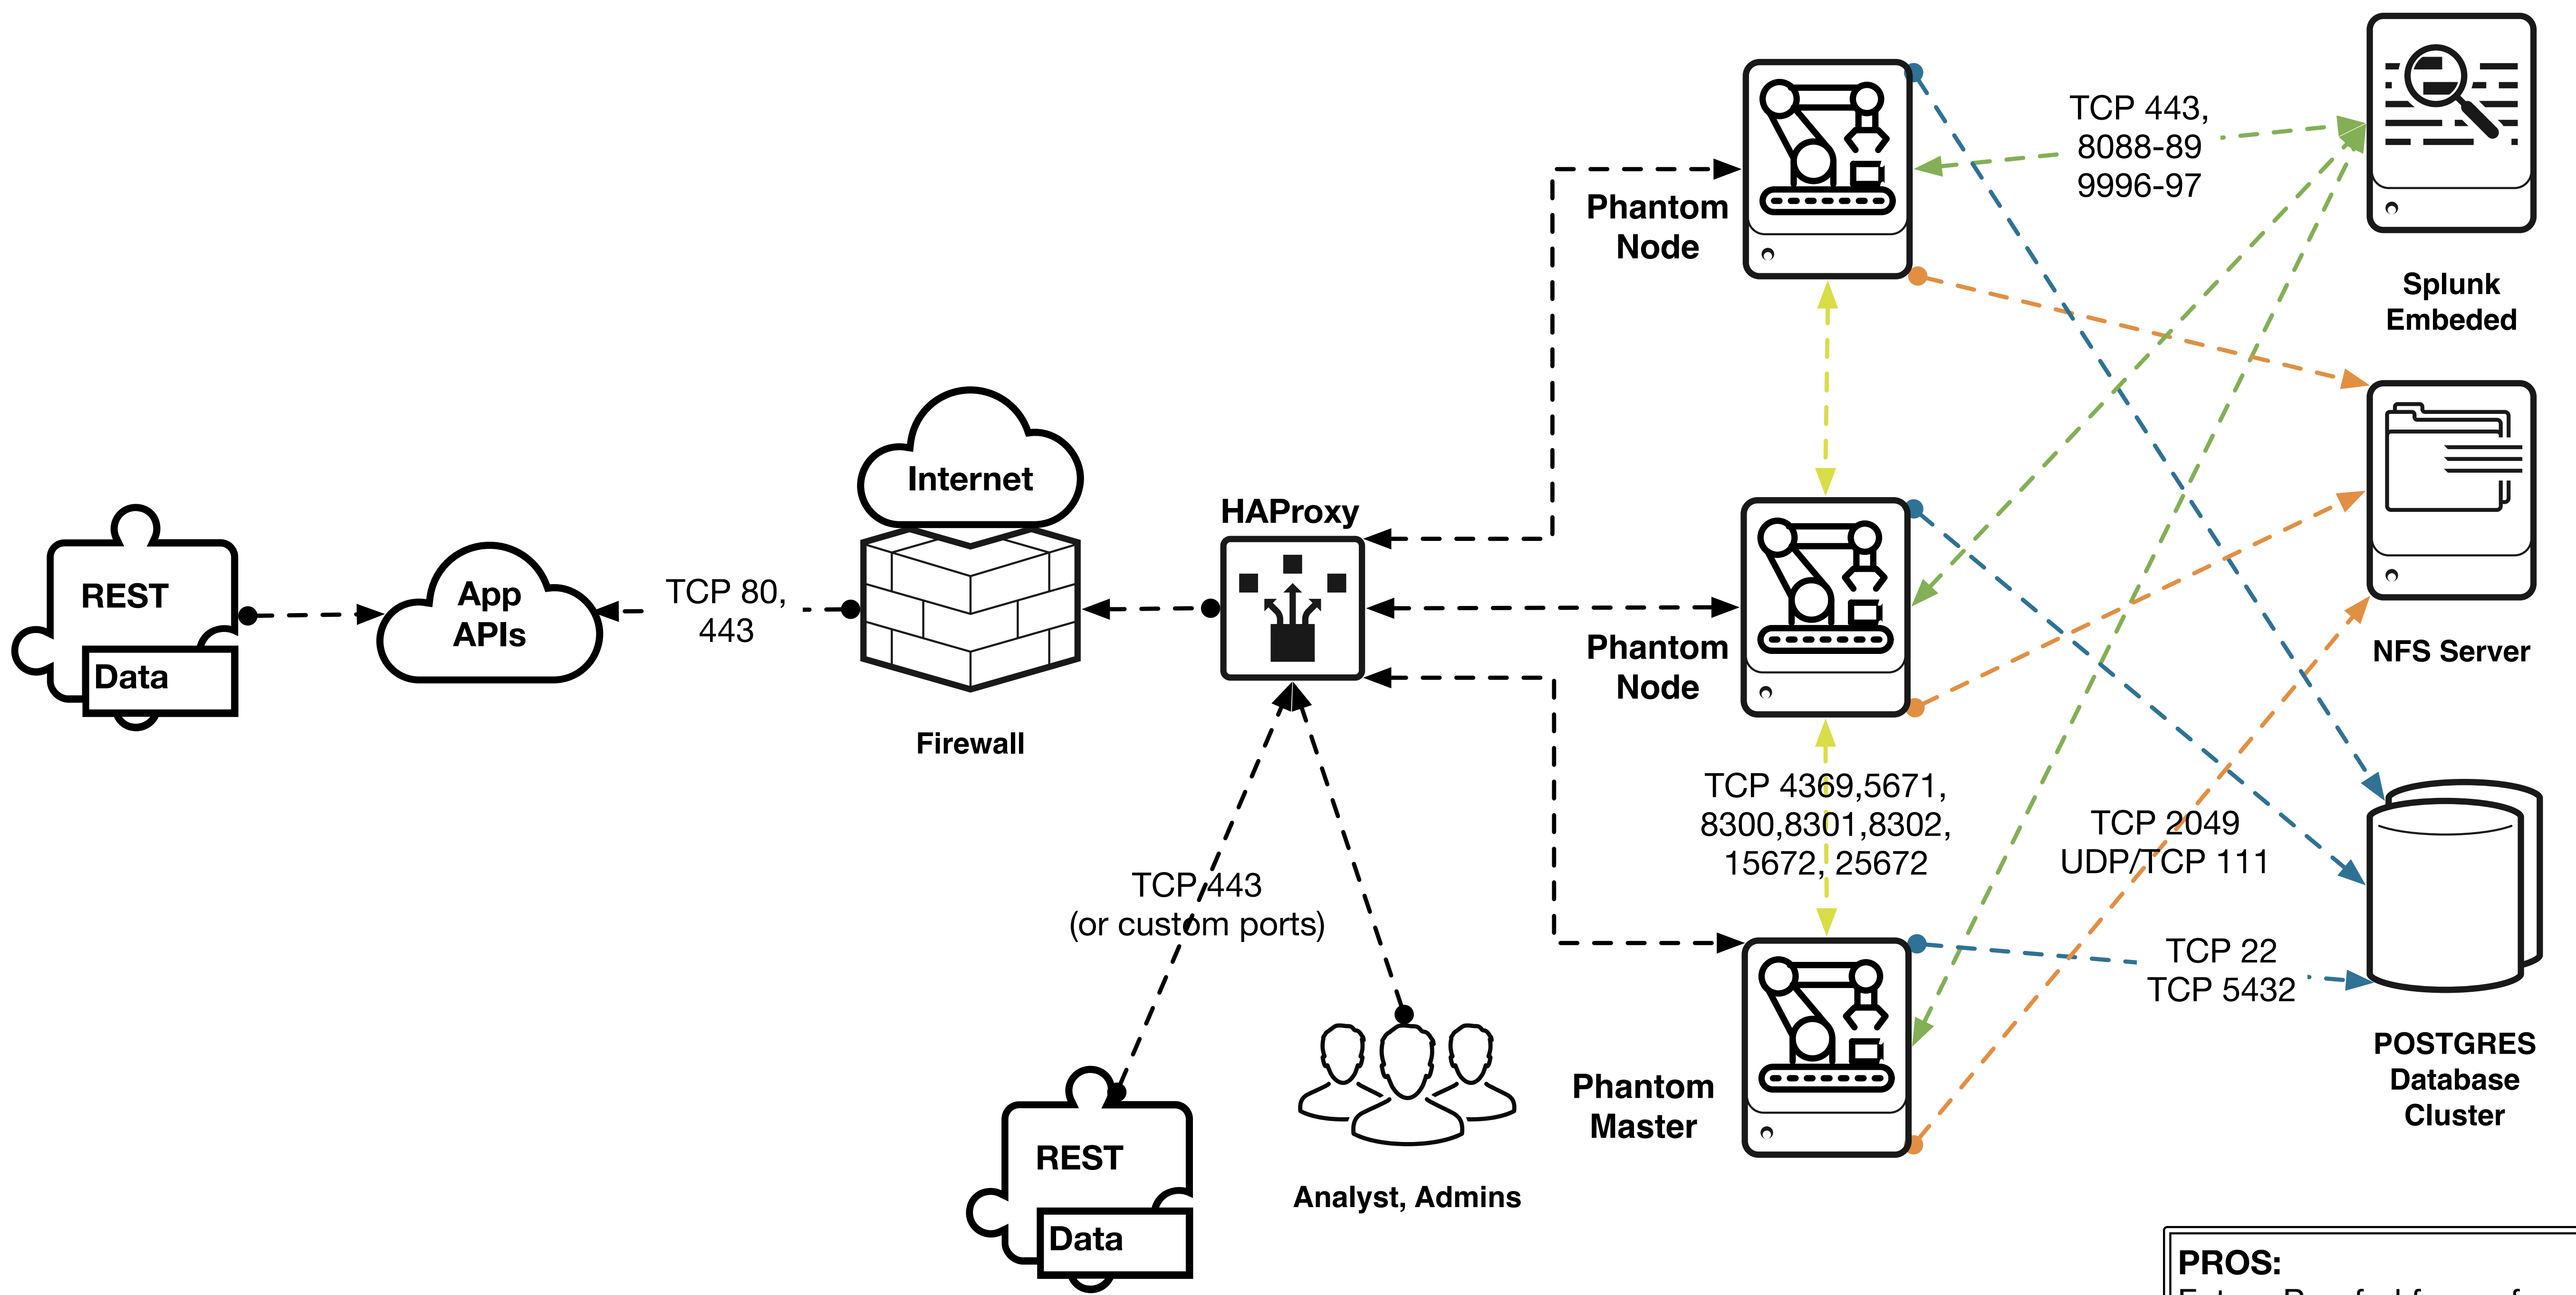

PROS:
Future Proofed for performance and no additional configuration required
Highest capacity design possible

CONS:
Complicated implementation
Requires a lot of unused resource until load is realized
Site to Site Sync isn't possible data is single homed to the regional infrastructure
Data is single homed to the sent site.

Production Node Drive Mappings	**Based on size (GB):**	**200**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	40.0
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	20.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	100.0
/dev/mapper/centos-tmp	/tmp	10.0
/dev/mapper/centos-var	/var	10.0
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	1.0
/dev/mapper/centos-var_tmp	/var/tmp	10.0
/dev/mapper/centos-home	/home	20.0
	Total:	**211.0**

Production Database Drive Mappings	**Based on size (GB):**	**1536**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	307.2
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	0.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	460.8
/dev/mapper/centos-tmp	/tmp	76.8
/dev/mapper/centos-var	/var	76.8
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	7.7
/dev/mapper/centos-var_tmp	/var/tmp	76.8
/dev/mapper/centos-home	/home	153.6
	Total:	**1159.7**

Production File Drive Mappings	**Based on size (GB):**	**1536**
Device	Mountpoint	Size (GB)
<glusterfs_hostname>:/apps	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/app_states	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/scm	/<phantom_install_dir>/local_data/app_states	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/scm	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/tmp/shared	153.6
<glusterfs_hostname>:/vault	/<phantom_install_dir>/vault	768.0
	Total:	**1536.0**

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 443	Use for Sending Notables to Phantom
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

POSTGRES Instance

Port	Purpose
TCP 22	Used for administering the OS that POSTGRES is running on
TCP 5432	PostgreSQL Service. Can be blocked if the DB server is a different host than the shared services node.

Cluster Message Queue

Port	Purpose
TCP 4369	RabbitMQ / Erlang port mapper. All cluster nodes must be able to talk to each-other on this port.
TCP 5671	RabbitMQ service. All cluster nodes must be able to talk to each-other on this port.
TCP 8300	Consul RPC services. All cluster nodes must be able to talk to each-other on this port.
TCP 8301	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 8302	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 15672	RabbitMQ admin UI and HTTP api service. UI is disabled by default. All cluster nodes must be able to talk to each-other on this port.
TCP 25672	RabbitMQ internode communications. All cluster nodes must be able to talk to each-other on this port.

Cluster Node

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments:

REVISIONS			
REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

C1E - High Capacity Clustered Deployment with Splunk Enterprise Integration

Intentionally left blank for diagram below

Splunk Phantom Cluster Sizing Plan

Cluster Type	Workloads with active playbooks	DB/ Common Node CPU cores	DB/Common Node Memory GB	Phantom Node CPU cores	Phantom Node Memory GB	Number of Phantom Nodes
XLarge	>50,000	32	64	16-32	32	8
Large	Up to 25,000-50,000 per hour	16	64	8	16	8
Medium	Up to 25,000 events per hour	16	32	8	16	5
Small	Up to 10,000 events per hour	8	32	4	8	3

Recommended Sizing

SpVA: C1E

splunk> phantom

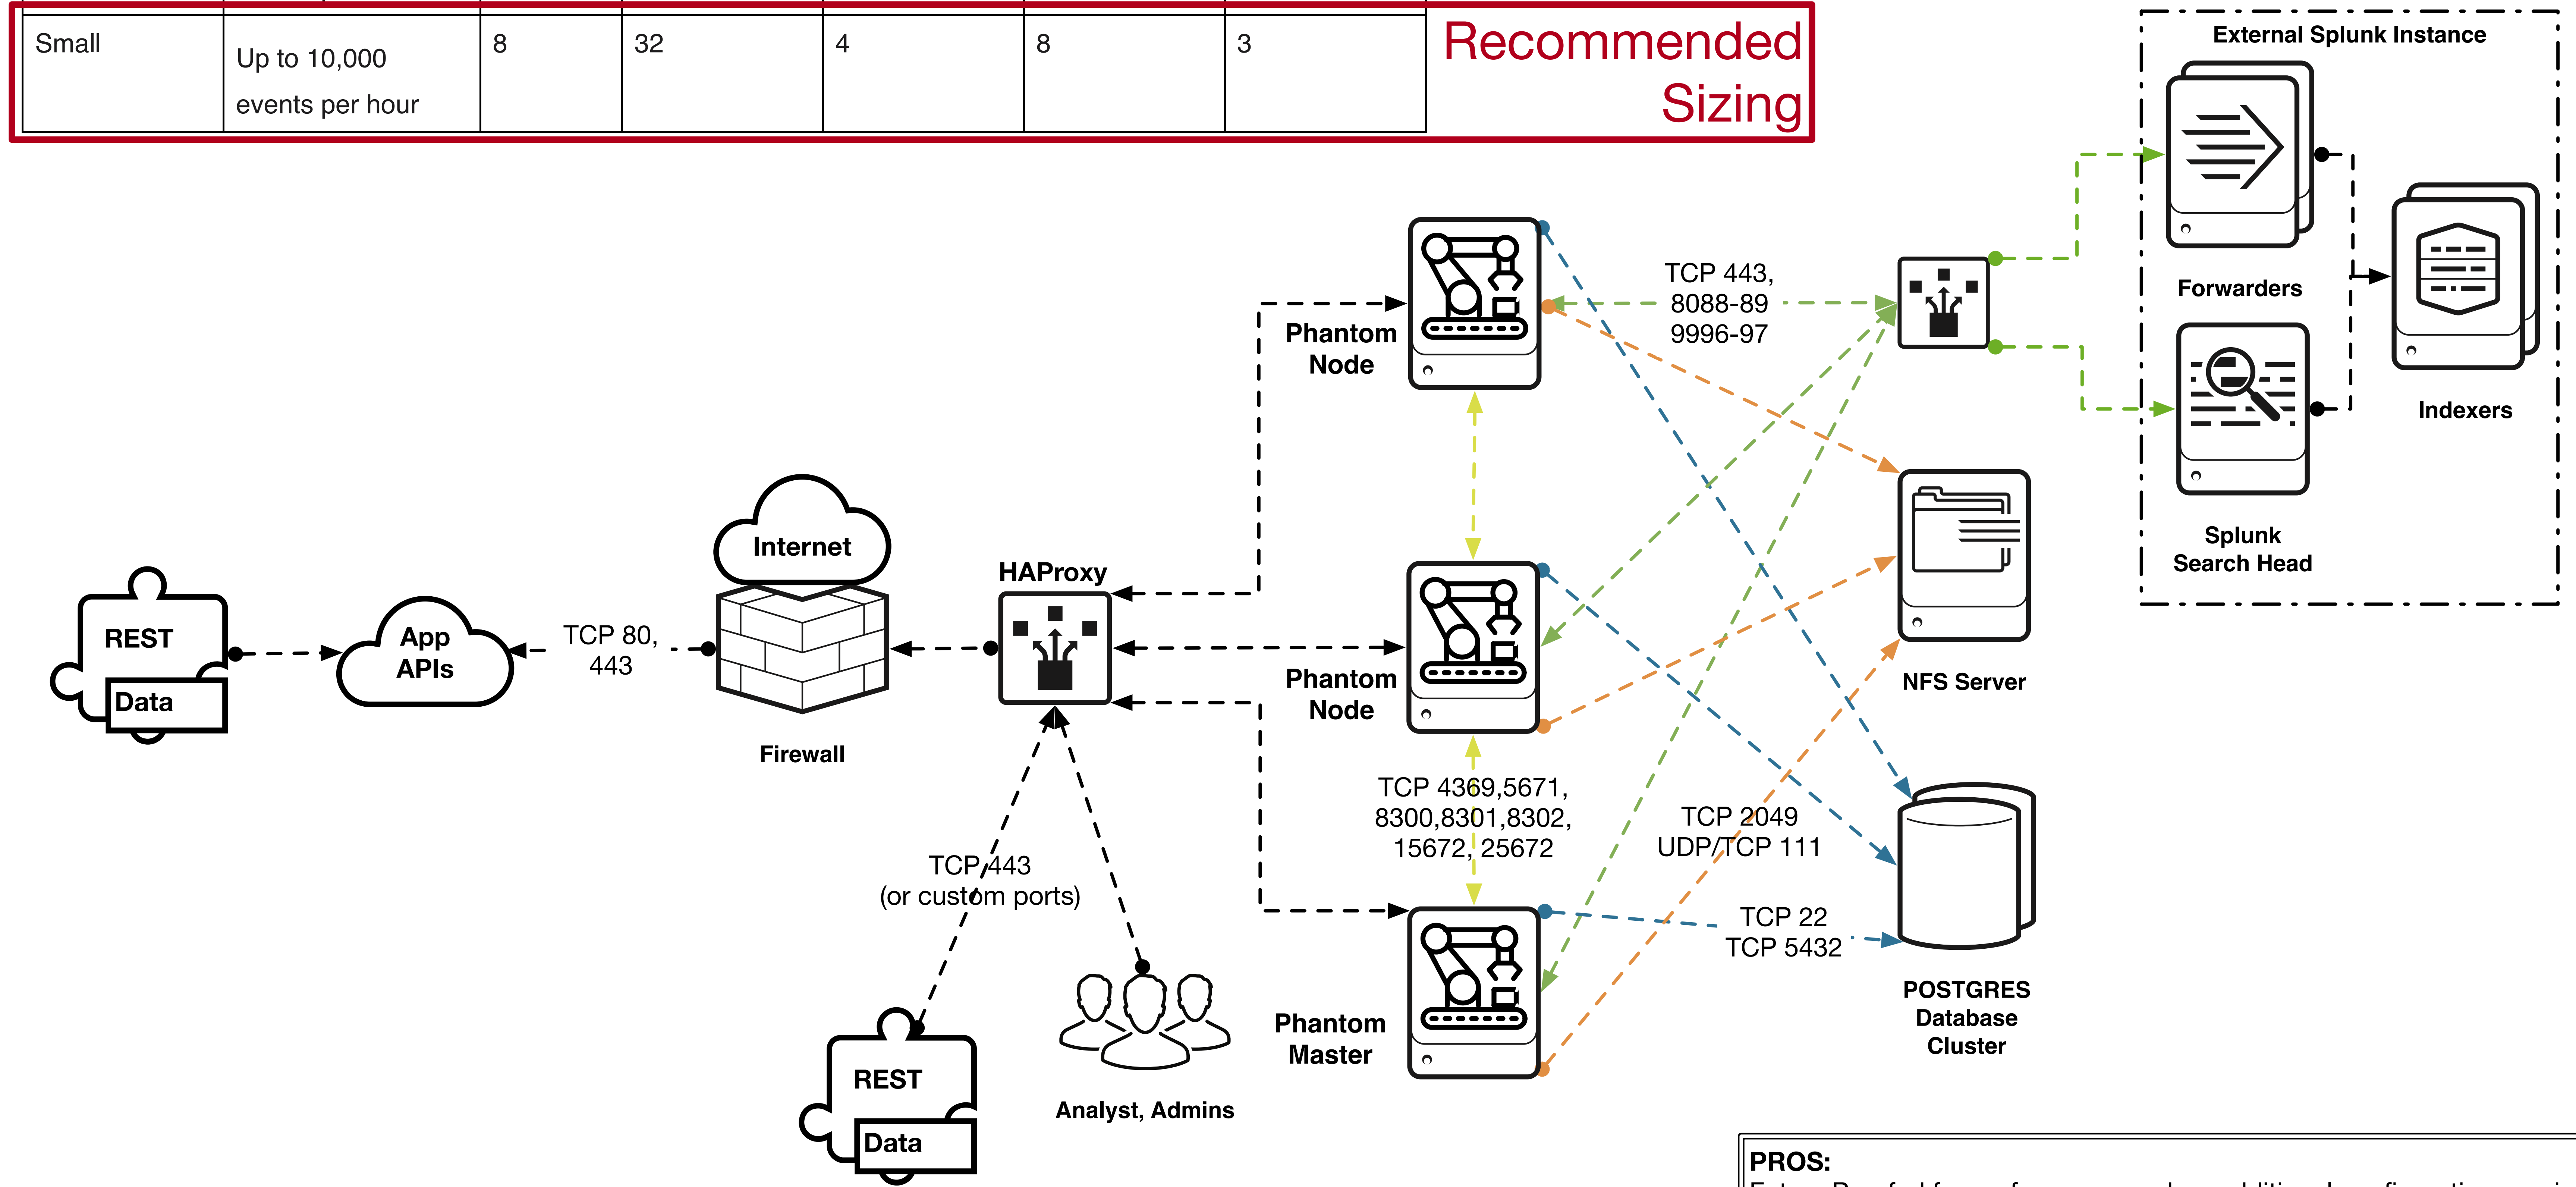

Production Node Drive Mappings	**Based on size (GB):**	**200**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	40.0
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	20.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	100.0
/dev/mapper/centos-tmp	/tmp	10.0
/dev/mapper/centos-var	/var	10.0
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	1.0
/dev/mapper/centos-var_tmp	/var/tmp	10.0
/dev/mapper/centos-home	/home	20.0
	Total:	**211.0**

Production Database Drive Mappings	**Based on size (GB):**	**1536**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	307.2
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	0.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	460.8
/dev/mapper/centos-tmp	/tmp	76.8
/dev/mapper/centos-var	/var	76.8
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	7.7
/dev/mapper/centos-var_tmp	/var/tmp	76.8
/dev/mapper/centos-home	/home	153.6
	Total:	**1159.7**

Production File Drive Mappings	**Based on size (GB):**	**1536**
Device	Mountpoint	Size (GB)
<glusterfs_hostname>:/apps	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/app_states	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/scm	/<phantom_install_dir>/local_data/app_states	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/scm	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/tmp/shared	153.6
<glusterfs_hostname>:/vault	/<phantom_install_dir>/vault	768.0
	Total:	**1536.0**

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 443	Use for Sending Notables to Phantom
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

POSTGRES Instance

Port	Purpose
TCP 22	Used for administering the OS that POSTGRES is running on
TCP 5432	PostgreSQL Service. Can be blocked if the DB server is a different host than the shared services node.

Cluster Message Queue

Port	Purpose
TCP 4369	RabbitMQ / Erlang port mapper. All cluster nodes must be able to talk to each-other on this port.
TCP 5671	RabbitMQ service. All cluster nodes must be able to talk to each-other on this port.
TCP 8300	Consul RPC services. All cluster nodes must be able to talk to each-other on this port.
TCP 8301	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 8302	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 15672	RabbitMQ admin UI and HTTP api service. UI is disabled by default. All cluster nodes must be able to talk to each-other on this port.
TCP 25672	RabbitMQ internode communications. All cluster nodes must be able to talk to each-other on this port.

Cluster Node

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

PROS:
Future Proofed for performance and no additional configuration required
Highest capacity design possible

CONS:
Complicated implementation
Requires a lot of unused resource until load is realized
Site to Site Sync isn't possible data is single homed to the regional infrastructure
Data is single homed to the sent site.

Comments:

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

C1E+ - High Capacity Clustered Deployment with AWS Integrations or customer provided infrastructure

Intentionally left blank for diagram below

SpVA: C1E+

Splunk Phantom Cluster Sizing for Amazon Web Services

Cluster Type	Workloads with an active playbook	DB/ Common Node RDS size	Phantom Node AWS EC2 size	Number of Phantom Nodes	
XLarge	>50,000	db.m5.16xlarge	c5.4xlarge	8	
Large	Up to 25,000-50,000 per hour	db.m5.4xlarge	c5.2xlarge	8	
Medium	Up to 25,000 events per hour	db.r4.8xlarge	c5.2xlarge	5	
Small	Up to 10,000 events per hour	db.m5.2xlarge	c5.xlarge	3	Recommended Sizing

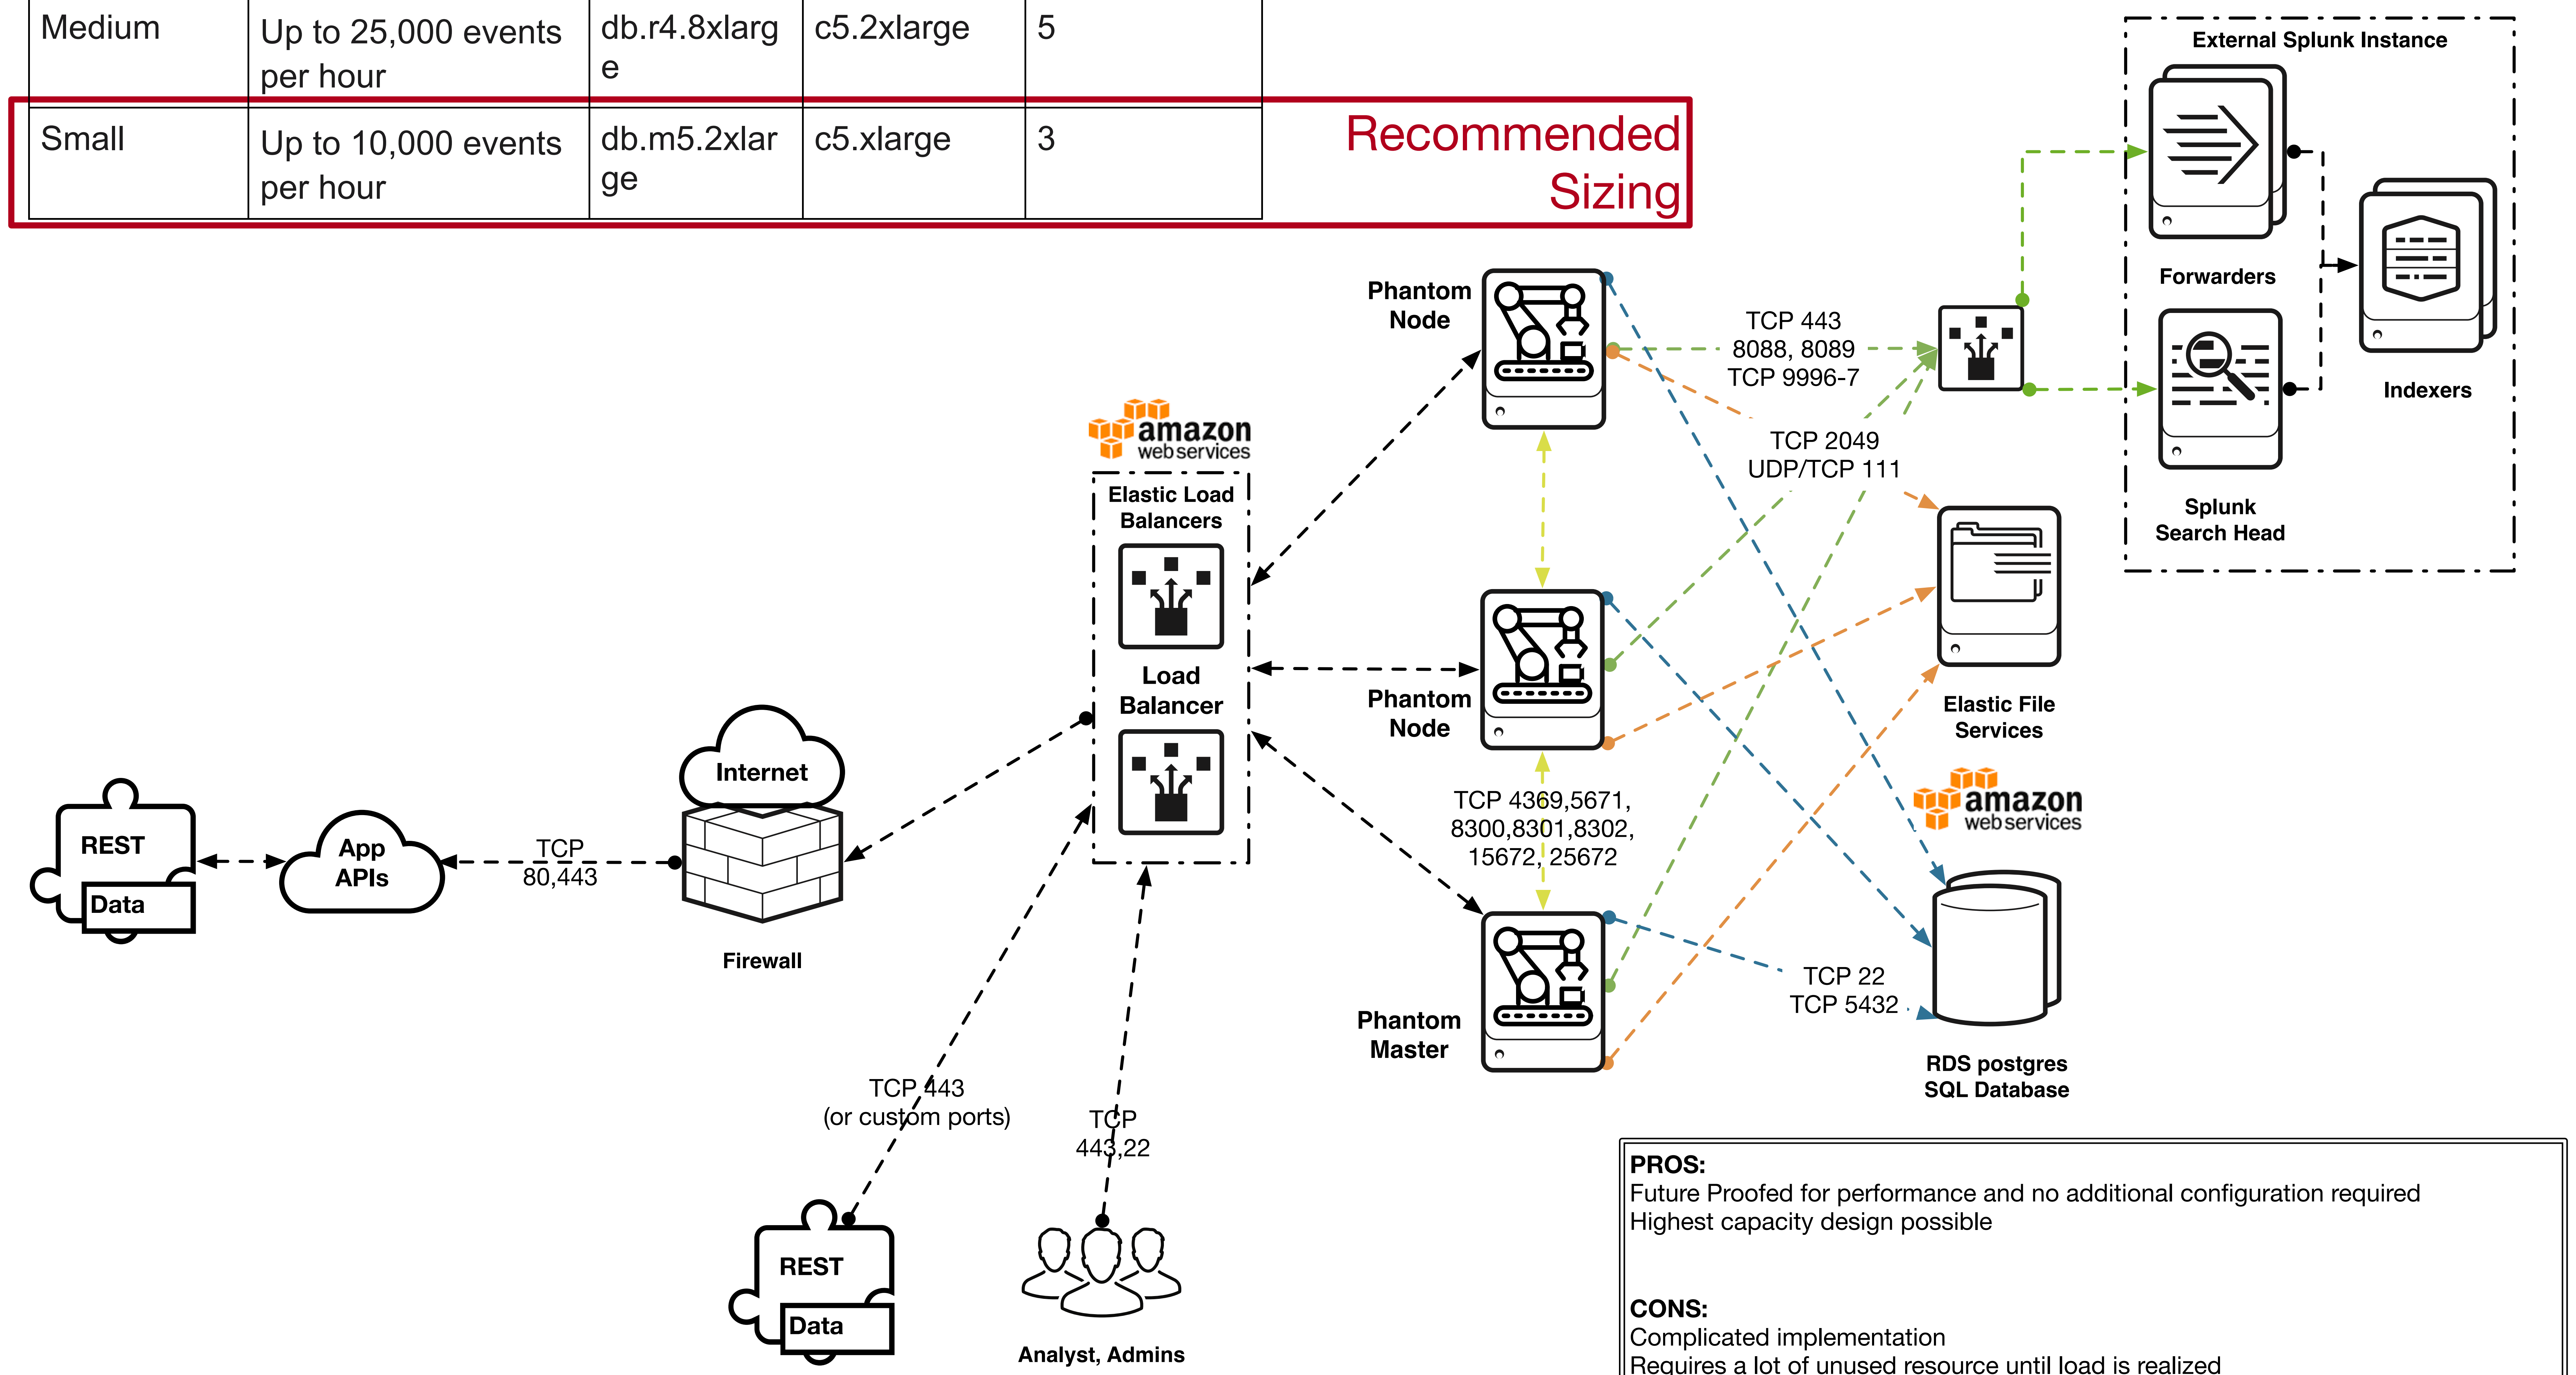

PROS:
Future Proofed for performance and no additional configuration required
Highest capacity design possible

CONS:
Complicated implementation
Requires a lot of unused resource until load is realized
Site to Site Sync isn't possible data is single homed to the regional infrastructure
Data is single homed to the sent site.

Production Node Drive Mappings	Based on size (GB):	200
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	40.0
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	20.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	100.0
/dev/mapper/centos-tmp	/tmp	10.0
/dev/mapper/centos-var	/var	10.0
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	1.0
/dev/mapper/centos-var_tmp	/var/tmp	10.0
/dev/mapper/centos-home	/home	20.0
	Total:	**211.0**

Production Database Drive Mappings	Based on size (GB):	1536
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	307.2
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	0.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	460.8
/dev/mapper/centos-tmp	/tmp	76.8
/dev/mapper/centos-var	/var	76.8
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	7.7
/dev/mapper/centos-var_tmp	/var/tmp	76.8
/dev/mapper/centos-home	/home	153.6
	Total:	**1159.7**

Production File Drive Mappings	Based on size (GB):	1536
Device	Mountpoint	Size (GB)
<glusterfs_hostname>:/apps	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/app_states	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/scm	/<phantom_install_dir>/local_data/app_states	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/scm	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/tmp/shared	153.6
<glusterfs_hostname>:/vault	/<phantom_install_dir>/vault	768.0
	Total:	**1536.0**

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 443	Use for Sending Notables to Phantom
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

POSTGRES Instance

Port	Purpose
TCP 22	Used for administering the OS that POSTGRES is running on
TCP 5432	PostgreSQL Service. Can be blocked if the DB server is a different host than the shared services node.

Cluster Message Queue

Port	Purpose
TCP 4369	RabbitMQ / Erlang port mapper. All cluster nodes must be able to talk to each-other on this port.
TCP 5671	RabbitMQ service. All cluster nodes must be able to talk to each-other on this port.
TCP 8300	Consul RPC services. All cluster nodes must be able to talk to each-other on this port.
TCP 8301	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 8302	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 15672	RabbitMQ admin UI and HTTP api service. UI is disabled by default. All cluster nodes must be able to talk to each-other on this port.
TCP 25672	RabbitMQ internode communications. All cluster nodes must be able to talk to each-other on this port.

Cluster Node

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

SpVA: C1E+

Splunk Phantom Cluster Sizing Plan

Cluster Type	Workloads with active playbooks	DB/ Common Node CPU cores	DB/Common Node Memory GB	Phantom Node CPU cores	Phantom Node Memory GB	Number of Phantom Nodes
XLarge	>50,000	32	64	16-32	32	8
Large	Up to 25,000-50,000 per hour	16	64	8	16	8
Medium	Up to 25,000 events per hour	16	32	8	16	5
Small	Up to 10,000 events per hour	8	32	4	8	3

Recommended Sizing

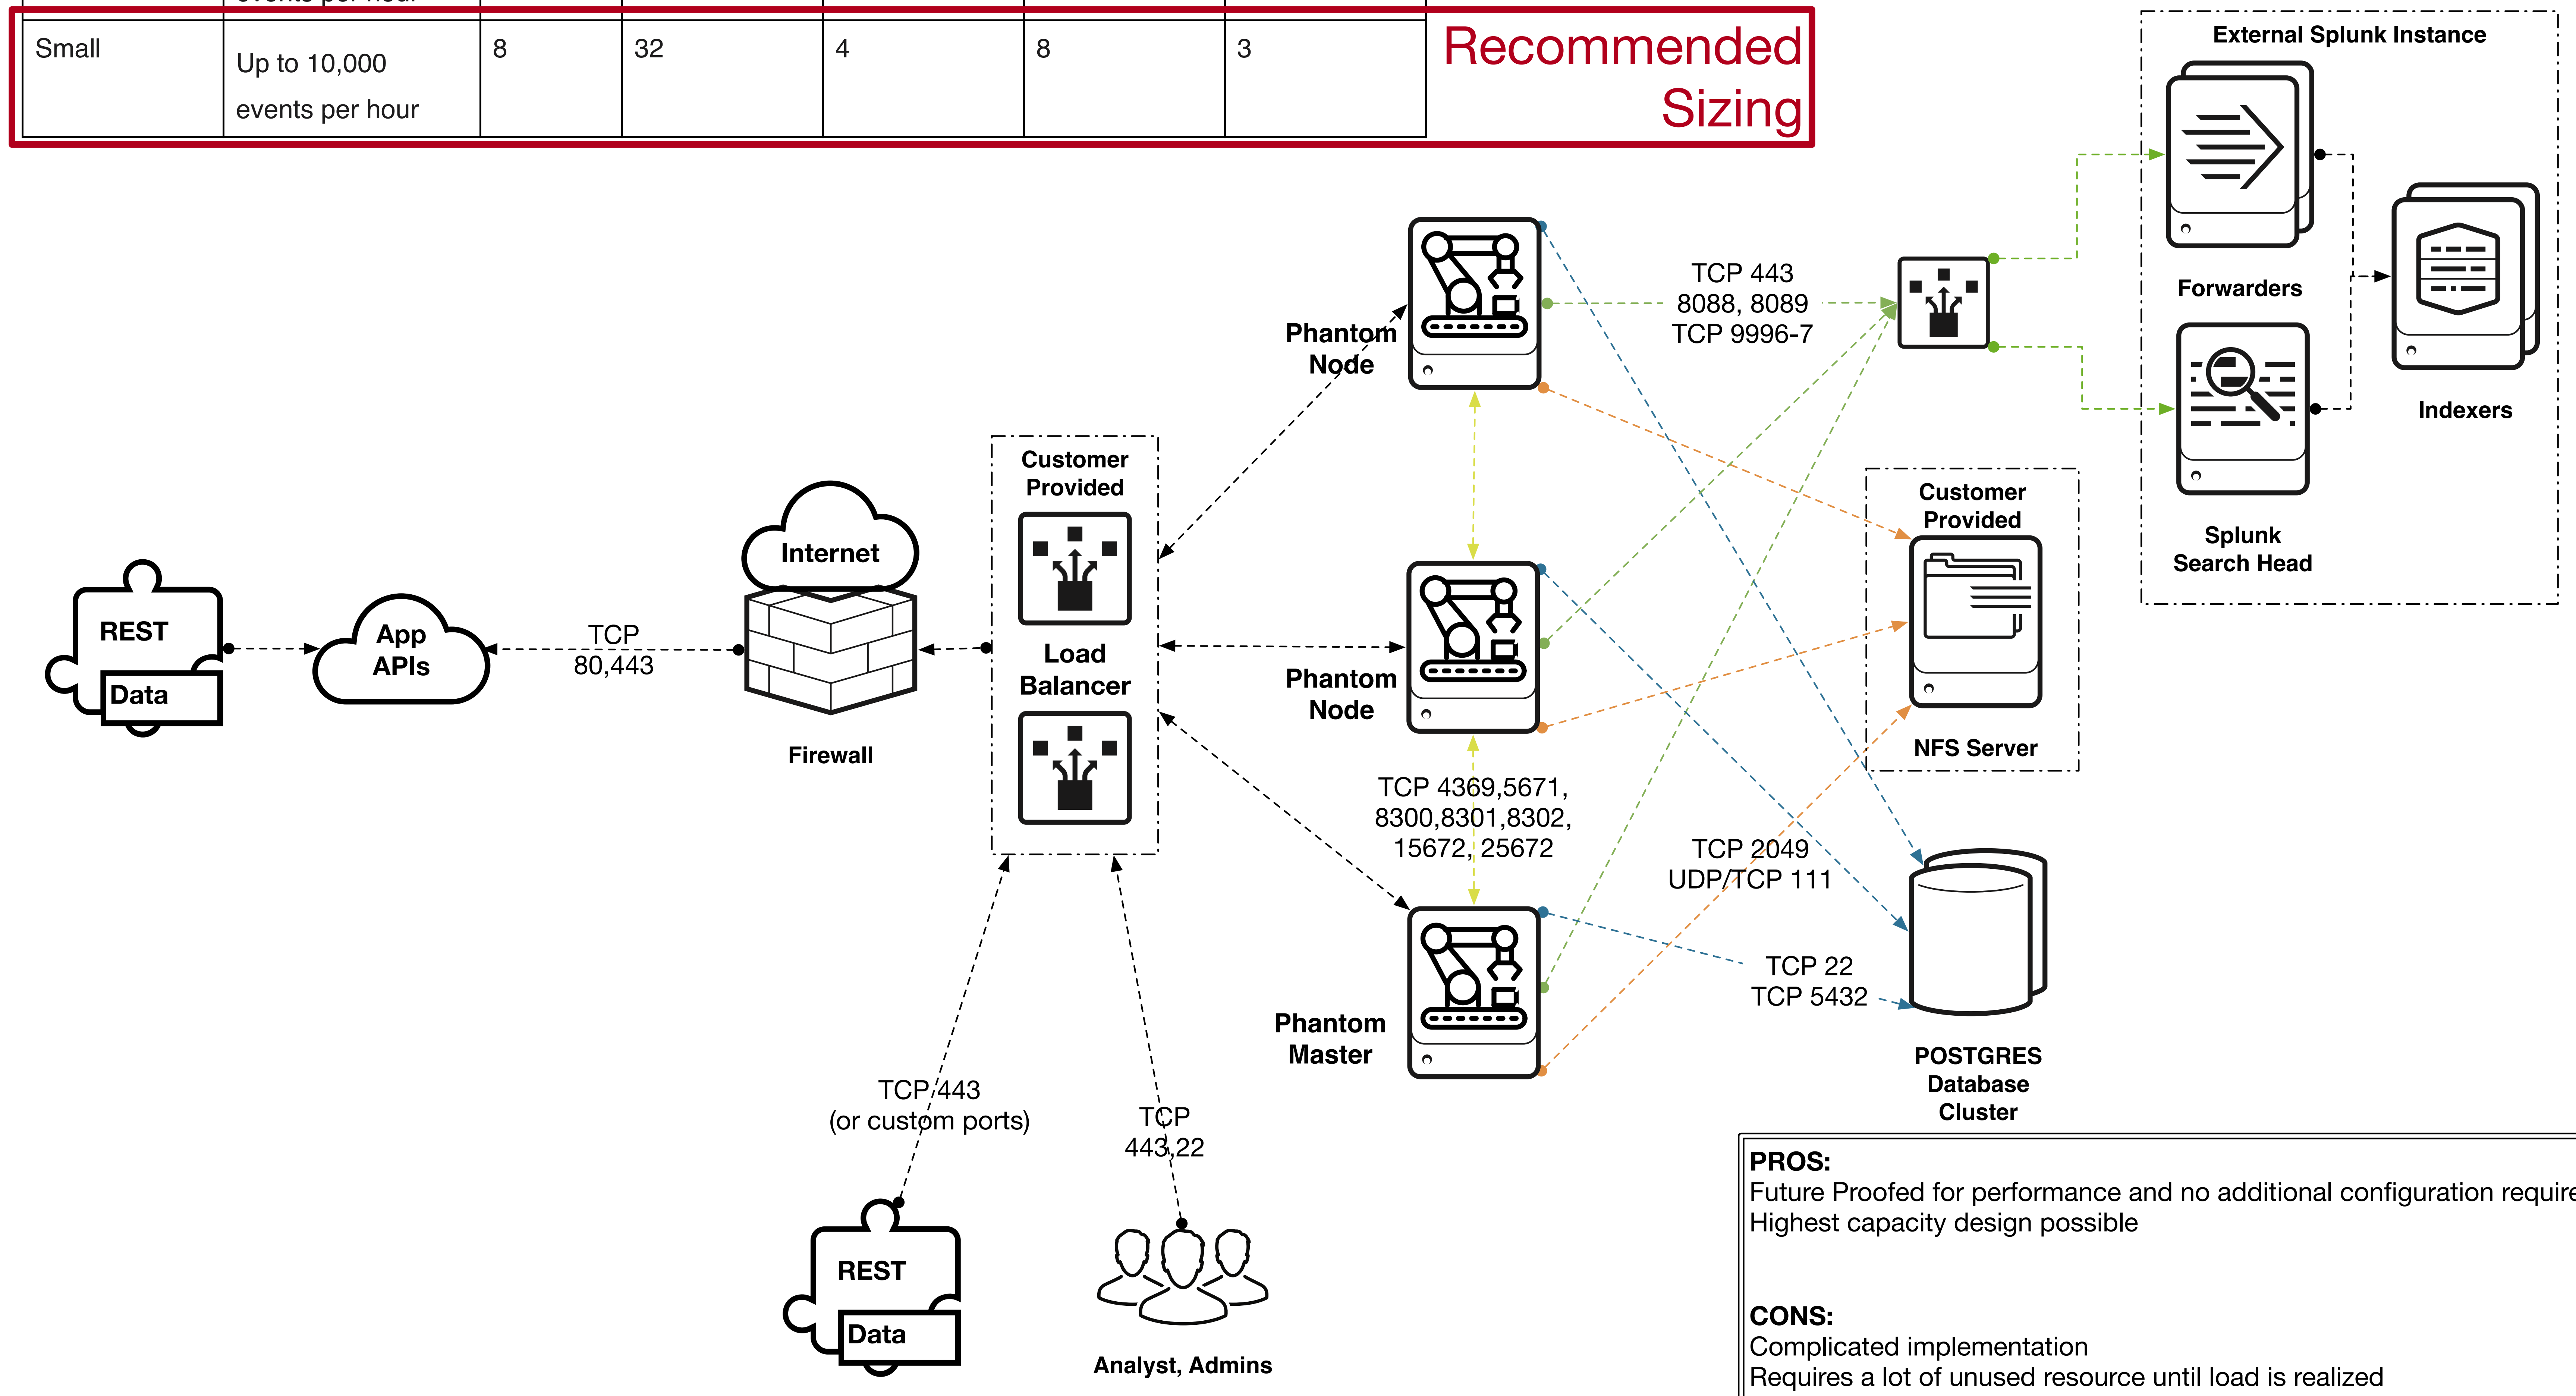

PROS:
Future Proofed for performance and no additional configuration required
Highest capacity design possible

CONS:
Complicated implementation
Requires a lot of unused resource until load is realized
Site to Site Sync isn't possible data is single homed to the regional infrastructure
Data is single homed to the sent site.

Production Node Drive Mappings	**Based on size (GB):**	**200**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	40.0
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	20.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	100.0
/dev/mapper/centos-tmp	/tmp	10.0
/dev/mapper/centos-var	/var	10.0
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	1.0
/dev/mapper/centos-var_tmp	/var/tmp	10.0
/dev/mapper/centos-home	/home	20.0
	Total:	**211.0**

Production Database Drive Mappings	**Based on size (GB):**	**1536**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	307.2
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	0.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	460.8
/dev/mapper/centos-tmp	/tmp	76.8
/dev/mapper/centos-var	/var	76.8
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	7.7
/dev/mapper/centos-var_tmp	/var/tmp	76.8
/dev/mapper/centos-home	/home	153.6
	Total:	**1159.7**

Production File Drive Mappings	**Based on size (GB):**	**1536**
Device	Mountpoint	Size (GB)
<glusterfs_hostname>:/apps	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/app_states	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/scm	/<phantom_install_dir>/local_data/app_states	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/scm	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/tmp/shared	153.6
<glusterfs_hostname>:/vault	/<phantom_install_dir>/vault	768.0
	Total:	**1536.0**

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 443	Use for Sending Notables to Phantom
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

POSTGRES Instance

Port	Purpose
TCP 22	Used for administering the OS that POSTGRES is running on
TCP 5432	PostgreSQL Service. Can be blocked if the DB server is a different host than the shared services node.

Cluster Message Queue

Port	Purpose
TCP 4369	RabbitMQ / Erlang port mapper. All cluster nodes must be able to talk to each-other on this port.
TCP 5671	RabbitMQ service. All cluster nodes must be able to talk to each-other on this port.
TCP 8300	Consul RPC services. All cluster nodes must be able to talk to each-other on this port.
TCP 8301	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 8302	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 15672	RabbitMQ admin UI and HTTP api service. UI is disabled by default. All cluster nodes must be able to talk to each-other on this port.
TCP 25672	RabbitMQ internode communications. All cluster nodes must be able to talk to each-other on this port.

Cluster Node

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

C1CE - High Capacity Clustered Deployment with Splunk Cloud

Intentionally left blank for diagram below

SpVA: C1CE

splunk> phantom

Splunk Phantom Cluster Sizing Plan

Cluster Type	Workloads with active playbooks	DB/ Common Node CPU cores	DB/Common Node Memory GB	Phantom Node CPU cores	Phantom Node Memory GB	Number of Phantom Nodes
XLarge	>50,000	32	64	16-32	32	8
Large	Up to 25,000-50,000 per hour	16	64	8	16	8
Medium	Up to 25,000 events per hour	16	32	8	16	5
Small	Up to 10,000 events per hour	8	32	4	8	3

Recommended Sizing

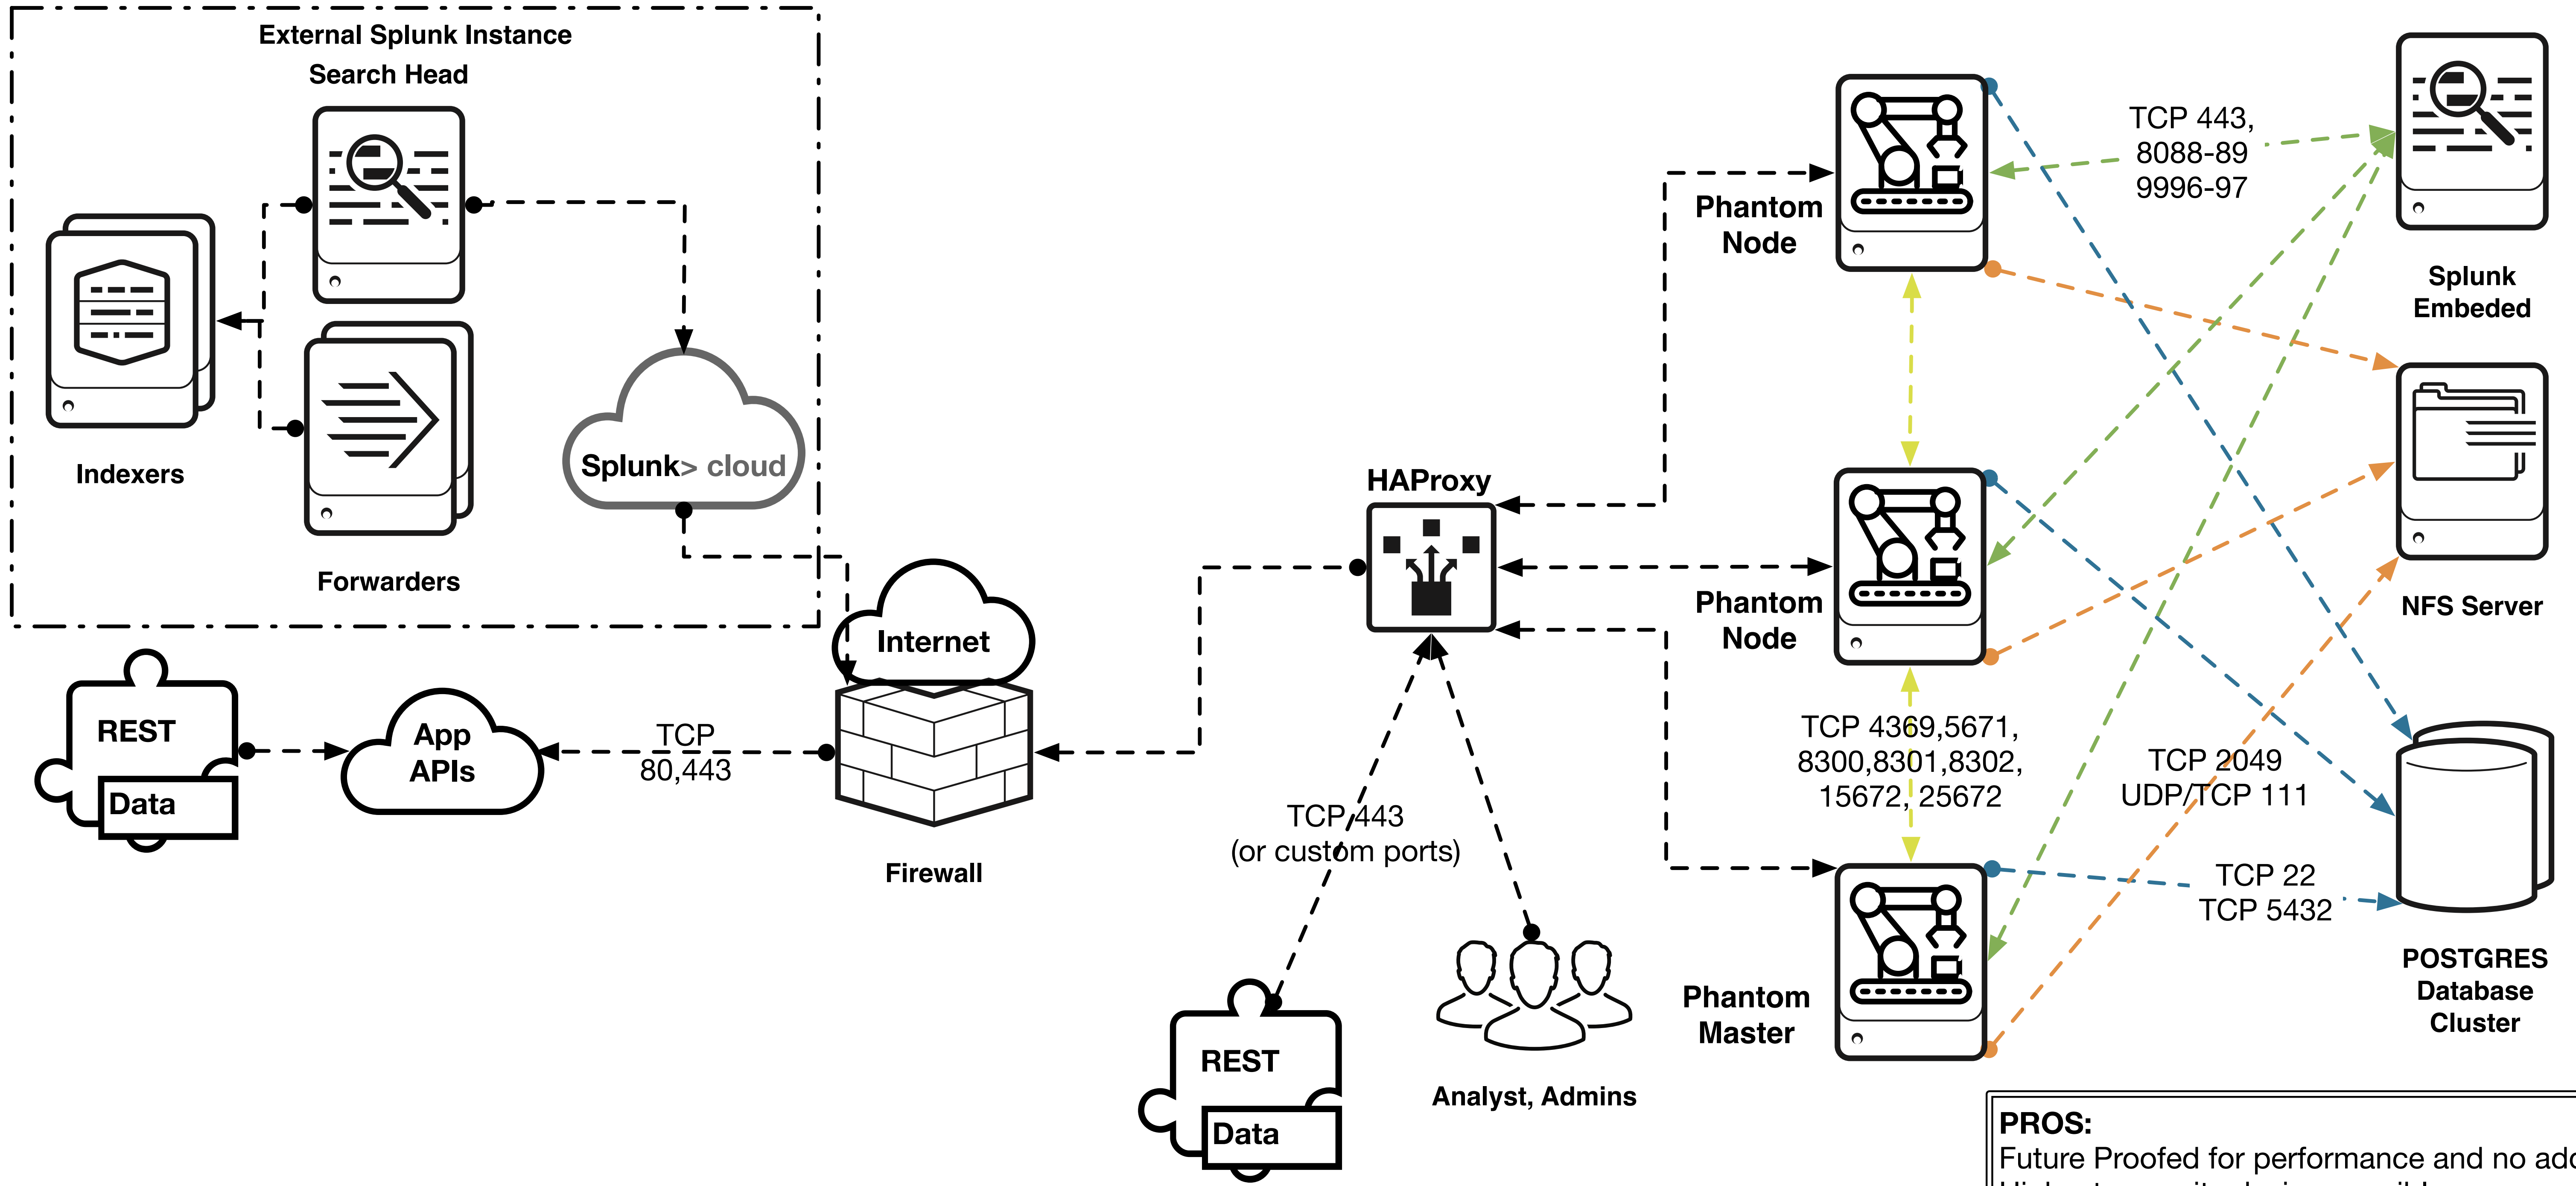

PROS:
Future Proofed for performance and no additional configuration required
Highest capacity design possible

CONS:
Complicated implementation
Requires a lot of unused resource until load is realized
Site to Site Sync isn't possible data is single homed to the regional infrastructure
Data is single homed to the sent site.

Production Node Drive Mappings	**Based on size (GB):**	**200**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	40.0
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	20.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	100.0
/dev/mapper/centos-tmp	/tmp	10.0
/dev/mapper/centos-var	/var	10.0
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	1.0
/dev/mapper/centos-var_tmp	/var/tmp	10.0
/dev/mapper/centos-home	/home	20.0
	Total:	**211.0**

Production Database Drive Mappings	**Based on size (GB):**	**1536**
Device	Mountpoint	Size (GB)
/dev/mapper/centos-root	/	307.2
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*	0.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*	460.8
/dev/mapper/centos-tmp	/tmp	76.8
/dev/mapper/centos-var	/var	76.8
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore	7.7
/dev/mapper/centos-var_tmp	/var/tmp	76.8
/dev/mapper/centos-home	/home	153.6
	Total:	**1159.7**

Production File Drive Mappings	**Based on size (GB):**	**1536**
Device	Mountpoint	Size (GB)
<glusterfs_hostname>:/apps	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/app_states	/<phantom_install_dir>/apps	153.6
<glusterfs_hostname>:/scm	/<phantom_install_dir>/local_data/app_states	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/scm	153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/tmp/shared	153.6
<glusterfs_hostname>:/vault	/<phantom_install_dir>/vault	768.0
	Total:	**1536.0**

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 443	Use for Sending Notables to Phantom
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

POSTGRES Instance

Port	Purpose
TCP 22	Used for administering the OS that POSTGRES is running on
TCP 5432	PostgreSQL Service. Can be blocked if the DB server is a different host than the shared services node.

Cluster Message Queue

Port	Purpose
TCP 4369	RabbitMQ / Erlang port mapper. All cluster nodes must be able to talk to each-other on this port.
TCP 5671	RabbitMQ service. All cluster nodes must be able to talk to each-other on this port.
TCP 8300	Consul RPC services. All cluster nodes must be able to talk to each-other on this port.
TCP 8301	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 8302	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 15672	RabbitMQ admin UI and HTTP api service. UI is disabled by default. All cluster nodes must be able to talk to each-other on this port.
TCP 25672	RabbitMQ internode communications. All cluster nodes must be able to talk to each-other on this port.

Cluster Node

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

C1CE+ - High Capacity Clustered Deployment with Splunk Cloud and AWS Integrations

Intentionally left blank for diagram below

SpVA: C1CE+ splunk> phantom

Splunk Phantom Cluster Sizing for Amazon Web Services

Cluster Type	Workloads with an active playbook	DB/ Common Node RDS size	Phantom Node AWS EC2 size	Number of Phantom Nodes
XLarge	>50,000	db.m5.16xlarge	c5.4xlarge	8
Large	Up to 25,000-50,000 per hour	db.m5.4xlarge	c5.2xlarge	8
Medium	Up to 25,000 events per hour	db.r4.8xlarge	c5.2xlarge	5
Small	Up to 10,000 events per hour	db.m5.2xlarge	c5.xlarge	3

Recommended Sizing

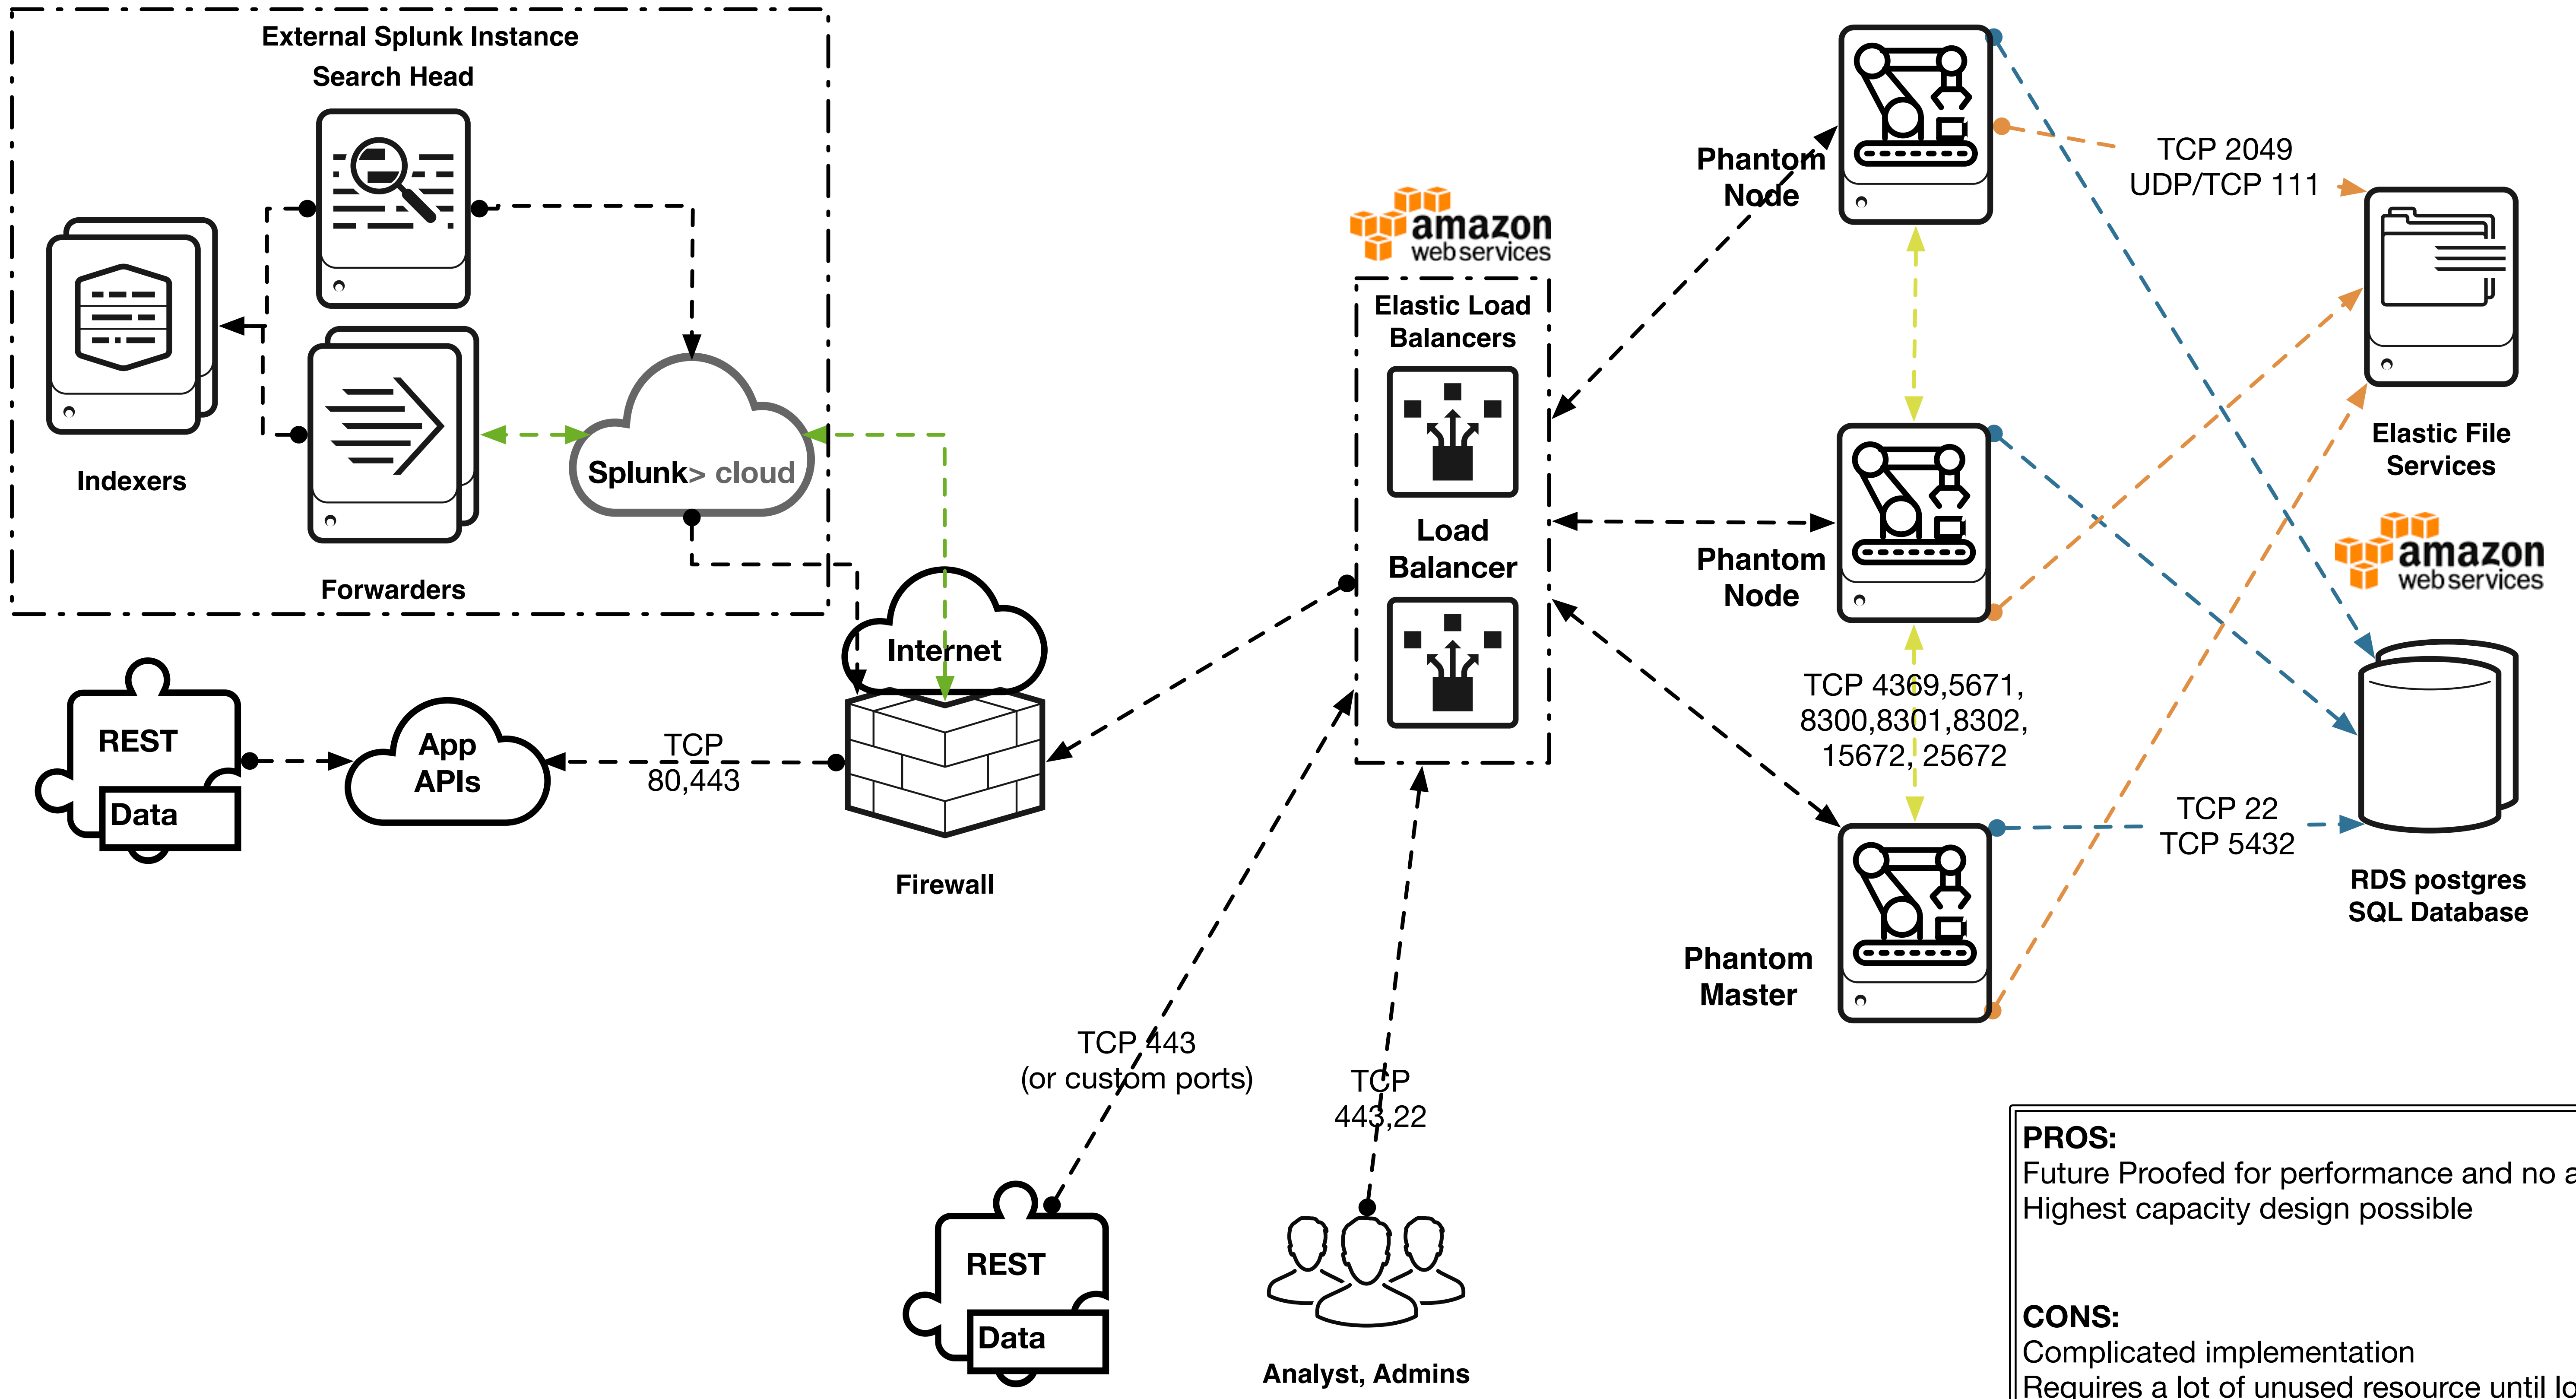

PROS:
Future Proofed for performance and no additional configuration required
Highest capacity design possible

CONS:
Complicated implementation
Requires a lot of unused resource until load is realized
Site to Site Sync isn't possible data is single homed to the regional infrastructure
Data is single homed to the sent site.

Production Node Drive Mappings		**Based on size (GB):**	**200**
Device	Mountpoint		Size (GB)
/dev/mapper/centos-root	/		40.0
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*		20.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*		100.0
/dev/mapper/centos-tmp	/tmp		10.0
/dev/mapper/centos-var	/var		10.0
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore		1.0
/dev/mapper/centos-var_tmp	/var/tmp		10.0
/dev/mapper/centos-home	/home		20.0
		Total:	**211.0**

Production Database Drive Mappings		**Based on size (GB):**	**1536**
Device	Mountpoint		Size (GB)
/dev/mapper/centos-root	/		307.2
/dev/mapper/centos-opt_phantom_vault	/opt/phantom/vault*		0.0
/dev/mapper/centos-opt_phantom_data	/opt/phantom/data*		460.8
/dev/mapper/centos-tmp	/tmp		76.8
/dev/mapper/centos-var	/var		76.8
/dev/mapper/centos-opt_phantom_keystore	/opt/phantom/keystore		7.7
/dev/mapper/centos-var_tmp	/var/tmp		76.8
/dev/mapper/centos-home	/home		153.6
		Total:	**1159.7**

Production File Drive Mappings		**Based on size (GB):**	**1536**
Device	Mountpoint		Size (GB)
<glusterfs_hostname>:/apps	/<phantom_install_dir>/apps		153.6
<glusterfs_hostname>:/app_states	/<phantom_install_dir>/apps		153.6
<glusterfs_hostname>:/scm	/<phantom_install_dir>/local_data/app_states		153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/scm		153.6
<glusterfs_hostname>:/tmp	/<phantom_install_dir>/tmp/shared		153.6
<glusterfs_hostname>:/vault	/<phantom_install_dir>/vault		768.0
		Total:	**1536.0**

NFS Server

Port	Purpose
TCP 2049	NFS Service
UDP 111	Portmapper service, needed for NFS
TCP 111	Portmapper service, needed for NFS

External Splunk Instance

Port	Purpose
TCP 443	Use for Sending Notables to Phantom
TCP 8088	Used as the HTTP Event Collector (HEC) and provides searching capabilities
TCP 8089	Used for the REST endpoint to send information to the Splunk Instances
TCP 9996-9997	Used for Universal Forwarder to either a forwarder or direct to the indexers

POSTGRES Instance

Port	Purpose
TCP 22	Used for administering the OS that POSTGRES is running on
TCP 5432	PostgreSQL Service. Can be blocked if the DB server is a different host than the shared services node.

Cluster Message Queue

Port	Purpose
TCP 4369	RabbitMQ / Erlang port mapper. All cluster nodes must be able to talk to each-other on this port.
TCP 5671	RabbitMQ service. All cluster nodes must be able to talk to each-other on this port.
TCP 8300	Consul RPC services. All cluster nodes must be able to talk to each-other on this port.
TCP 8301	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 8302	Consul internode communication. All cluster nodes must be able to talk to each-other on this port.
TCP 15672	RabbitMQ admin UI and HTTP api service. UI is disabled by default. All cluster nodes must be able to talk to each-other on this port.
TCP 25672	RabbitMQ internode communications. All cluster nodes must be able to talk to each-other on this port.

Cluster Node

Port	Purpose
TCP 22	Used for administering the OS that Phantom is running on. Can be limited to authorized administration networks, or blocked if you wish to use the OS console exclusively.
TCP 80	Convenience port for users who do not specify HTTPS when connecting to Phantom instance. Exists only to redirect connections to TCP 443. Can be blocked.
TCP 443	HTTPS interface for the web UI for Phantom, as well as REST access. Must be exposed to anything accessing the Phantom services.

Comments: This is our Default PS Recommendation

REVISIONS

REV	DESCRIPTION	DATE	APPROVED
1.0	Initial Build	6 Aug 18	

DRAWN BY	Architect
ISSUED TO	Approver

Company Name

M2E - High Capacity Clustered Deployment - Consult your Splunk Architect for a custom solution

Intentionally left blank for diagram below

M2CE – Consult your Splunk Architect for a custom solution

M2CE+ - Consult your Splunk Architect for a custom solution

1.7.2 Aligning Your Topology with Best Practices

You will need to keep your requirements and topology in mind in order to select the appropriate design principles and best practices for your deployment. Therefore, you should consider best practices only after you have completed Steps 1 and 2 of the Splunk SOAR Validated Architectures selection process above.

1.7.3 Best Practices: Tier-Specific Recommendations

Below you will find design principles and best practices recommendation for each deployment tier. Each design principle reinforces one or more of the SSVA pillars: Availability, Performance, Scalability, Security, and Manageability.

1.7.3.1 Automation and Case Management Tier Recommendations

DESIGN PRINCIPLES / BEST PRACTICES *(Your requirements will determine which practices apply to you)*		***SSVA PILLARS***				
		AVAILABILITY	***PERFORMANCE***	***SCALABILITY***	***SECURITY***	***MANAGEABILITY***
1	*Consider using SSDs for data volumes* *SSDs have reached economical prices and remove any possible IO limitations that are often the cause for unsatisfactory search performance.*		✓			
2	*Keep automation tier close (in network terms) to the user base* *Lowest possible network latency will have positive effect on user experience when using case management.*		✓			
3	*Use warm standby replication to minimize configuration and maintenance* *Warm standby ensures a copy of every event in the SOAR platform is protected against SOAR node failure.*	✓				
4	*Consider using LDAP/SAML auth whenever possible* *Centrally managing user identities for authentication purposes is a general enterprise best practice, simplifies management of your Splunk deployment and increases security.*			✓	✓	✓
5	*Ensure enough cores and memory to cover automation needs (start with 32GBs and 16 Cores, for a single instance)* *Every automation workflow requires Memory and CPU cores to execute. If there is memory pressure or no cores are available to run a playbook, the playbook will be queued, resulting in playbook and action delays for the user.*	✓	✓	✓		
6	*Avoid using multiple independent SOAR instances* *Independent SOAR instances do not allow sharing of Splunk artifacts created by users.*		✓	✓		✓
7	*Utilize git services for playbook replication between development and production automation tiers*		✓		✓	✓
8	*Monitor critical automation metrics*	✓	✓			

	Splunk provides you with a monitoring console that provides key performance metrics on how your automation tier is performing. This includes disk usage, CPU and memory utilization, as well as detailed metrics of internal Splunk components (processes, process utilization, and queues).					

Summary & Next Steps

This white paper provided a general introduction to Splunk SOAR Validated Architectures and ensures that your organization's requirements are being met in the most cost-effective, manageable, and scalable way possible. SSVAs offer best practices and design principles built upon the following foundational pillars:

- Availability
- Performance
- Scalability
- Security
- Manageability

This white paper has also covered the 3-step Splunk SOAR Validated Architectures selection process:

1) Definition of requirements,
2) Choosing a topology, and
3) Applying design principles and best practices.

Now that you are familiar with the multiple benefits of Splunk SOAR Validated Architectures, we hope you are ready to move forward with the process of choosing a suitable deployment topology for your organization.

1.8 Next Steps

So, what comes after choosing a Validated Architecture? The next steps on your journey to a working environment include:

Customizations

- Consider any necessary customizations your chosen topology may need to meet specific requirements.

Deployment Model

- Decide on deployment model (bare metal, virtual, cloud). We highly recommend virtual whether they are “on premise” or cloud based.

System

- Select your technology (servers, storage, operating systems) according to Splunk system requirements. (https://docs.splunk.com/Documentation/SOAR/4.8/Install/Requirements)

Sizing

- Gather all the relevant data you will need to size your deployment (data ingest, expected playbook volume, data retention needs, replication, etc.) Discuss these requirements with your assigned Splunk SOAR Security Solutions Architect.

Staffing

- Evaluate your staffing needs to implement and manage your deployment. This is an essential part of building out a Splunk Center of Excellence.

We are here to assist you throughout the Validated Architectures process and with next steps. Please feel free to engage your Splunk Account Team with any questions you might have. Your Account Team will have access to the full suite of technical and architecture resources within Splunk and will be happy to provide you with further information.

Happy Splunking!

Appendix

This section contains additional reference information used in the SSVAs.

1.9 Appendix "A": SSVA Pillars Explained

Pillar	Description	Primary Goals / Design Principles
Availability	*The ability to be continuously operational and able to recover from planned and unplanned outages or disruptions.*	1. *Eliminate single points of failure / Add redundancy* 2. *Detect planned and unplanned failures/outages* 3. *Tolerate planned/unplanned outages, ideally automatically* 4. *Plan for rolling upgrades*
Performance	*The ability to effectively use available resources to maintain optimal level of service under varying usage patterns.*	1. *Add hardware to improve performance; compute, storage, memory.* 2. *Eliminate bottlenecks 'from the bottom up'* 3. *Exploit all means of concurrent processing* 4. *Exploit locality (i.e. minimize distribution of components)* 5. *Optimize for the common case (80/20 rule)* 6. *Avoid unnecessary generality* 7. *Time shift computation (pre-compute, lazily compute, share/batch compute)* 8. *Trade certainty and accuracy for time (randomization, sampling)*
Scalability	*The ability to ensure that the system is designed to scale on all tiers and handle increased workloads effectively.*	1. *Scale vertically and horizontally* 2. *Separate functional components that need to be scaled individually* 3. *Minimize dependencies between components* 4. *Design for known future growth as early as possible* 5. *Introduce hierarchy in the overall system design*
Security	*The ability to ensure that the system is designed to protect data as well as configurations/assets while continuing to deliver value.*	1. *Design for a secure system from the start* 2. *Employ state-of-the art protocols for all communications* 3. *Allow for broad-level and granular access to event data* 4. *Employ centralized authentication* 5. *Implement auditing procedures* 6. *Reduce attack or malicious use surface area*

Manageability	*The ability to ensure the system is designed to be centrally operable and manageable across all tiers.*	*1. Provide a centralized management function* *2. Manage configuration object lifecycle (source control)* *3. Measure and monitor/profile application (Splunk) usage* *4. Measure and monitor system health*

1.10 Appendix "B": Topology Components

Tier	Component	Icon	Description	Notes
Automation and Case Management	*SOAR (PH) Node*	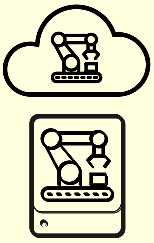	*A SOAR tenant or node provides the UI for Splunk users and provides case management and playbook development activities.*	*SOAR tenant is software as a service instance for your organization.* *SOAR Node is a dedicated Splunk SOAR appliance in distributed deployments.* *SOAR Node is frequently virtualized to provide vertical scalability and easy failure recovery, provided they are deployed with the appropriate CPU and memory resources.*
	SOAR Cluster	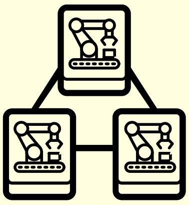	*A SOAR cluster provides the UI for Splunk users and provides case management and playbook development activities.*	*SOAR clusters require dedicated servers of ideally identical system specifications.* *SOAR clusters are frequently virtualized, provided they are deployed with the appropriate CPU and memory resources.*
	Automation Broker	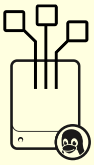	*A docker container that loads applications and allows SOAR to interact with on premise systems or services.*	*SOAR Automation Broker provides access to systems and resources for on premise or non-internet accessible resources.* *SOAR Automation Broker requires internet access to communicate to SOAR services.*
Shared Services	*Gluster File Server (gFS)*	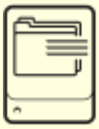	*Gluster file server provides an open-source secure file server.*	*Gluster file servers provide file storage activities for the case management capabilities or playbook automation.*
	External Database (Ext DB)	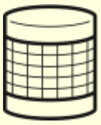	*External database is a PostgreSQL database or database server that provides the core UI data and storage for all the actions.*	*Externalizing the PostgreSQL database will improve SOAR performance but will keep the data to only one site. This database can be clustered also to improve data resiliency.*

			clustered environments.	
	HA Proxy		*HA Proxy is an open-source load balancer for the SOAR Clustered configurations.*	*HA Proxy is easily configured and maintains its node relationship automatically. However, it does have limited enterprise functionality. For robust cluster configurations, Splunk recommends appropriate enterprise network load balances for mission critical cluster operations.*
Reporting	*Search Head (SH)*		*The search head provides the UI for Splunk users and coordinates scheduled search activity.*	*Search heads are dedicated Splunk instances in distributed deployments. Search heads can be virtualized for easy failure recovery, provided they are deployed with appropriate CPU and memory resources.*
	Search Head Cluster (SHC)		*A search head cluster is a pool of at least three clustered Search Heads. It provides horizontal scalability for the search head tier and transparent user failover in case of outages.*	*Search head clusters require dedicated servers of ideally identical system specifications. Search head cluster members can be virtualized for easy failure recovery, provided they are deployed with appropriate CPU and memory resources.*