



Well-architecture Resilience Key Questions	Well-architecture Sub Categories for Resilience	Does Appranix help satisfy the requirement?	How does Appranix help with well-architecture requirements!
What reliability targets and metrics have you defined for your application? (Availability targets, such as Service Level Agreements (SLA) and Service Level Objectives (SLO), and Recovery targets, such as Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO), should be defined and tested to ensure application reliability aligns with business requirements.)	Recovery targets to identify how long the workload can be unavailable (Recovery Time Objective) and how much data is acceptable to lose during a disaster (Recovery Point Objective).	Yes	Appranix tracks and maintains RPO based on policies with local region and remote region protections. Users can adjust the policies per their organization RPO requirements and Appranix will ensure data backup, replication and lifecycle requirements based on the policies to deliver the required protection. Appranix recovers not only individual VMs but all other Azure resources to drastically reduce RTO for protected distributed applications running on Azure. Appranix continues to monitor the recovery times based on RTOs so an organizations can improve the recovery times as well.
	Availability targets such as Service Level Agreements (SLAs) and Service Level Objectives (SLOs)	Yes	By defining RPO and RTO per application Azure resources, also called as Appranix Cloud Assemblies, Site Reliability Engineers or Cloud Operations teams can further derive and monitor SLOs and SLAs for a particular cloud application.
	Availability metrics to measure and monitor availability such as Mean Time To Recover (MTTR) and Mean Time Between Failure (MTBF).	Yes	Appranix does not directly calculate and measure MTTR and MTBF. SREs and CloudOps team could use monitoring systems to measure these metrics, however, Appranix helps to identify the MTTR using internally built-in RTO metrics over time.

	Composite SLA for the workload derived using the Azure SLAs for all relevant resources	Yes	Yes, users can easily derive SLAs for composite workloads using Appranix's Cloud Assemblies. Cloud Assemblies are an automated and continuously updated set of Azure resources that roughly represents a composite cloud application.
	SLAs for all internal and external dependencies.	Yes	An Azure applications/composite workloads dependencies are automatically protected and recovered using Appranix's native capabilities. External dependencies recovery times can be controlled with pre- and post-webhooks mechanism.
	Independent availability and recovery targets for critical application subsystems and scenarios.Has the application been decomposed into distinct subsystems with independent availability and recovery targets?	Yes	Users can measure RPO and RTO targets for decomposed subsystems with separate Cloud Assemblies.
How have you ensured that your application architecture is resilient to failures? (Resilient application architectures should be designed to recover gracefully from failures in alignment with defined reliability targets.)	Deployed the application across multiple regions.Does the application support multi-region deployments for failover purposes?	Yes	Appranix helps achieve a level of multi-region deployment for customers as they may not have all the expertise and resources to deploy across multiple regions.
	Deployed the application across Availability Zones within a region.Is the application designed to use Availability Zones within a region?	Partially	Assuming that Azure users have appropriately designed their applications to run on multiple Availability zones, Appranix provides further protection and recovery capabilities across another region with the same level of Availability Zone replication and deployments. For example, if users have deployed their applications across 3 zones in a region when Appranix fails over their application to another region, all their Availability Zone configurations will be retained and Appranix automatically distributes those applications across as many zones as the primary region. In cases where the number of Availability Zones are less in certain regions, Appranix automatically reconfigures to only the available number of zones to make sure the failed over applications can run optimally until the primary/production region to ready for failback.
	Performed Failure Mode Analysis (FMA) to identify fault-points and fault-modes.	Yes	Appranix helps with creating a "production-twin" for customers to do failure mode analysis. It is difficult for many Azure customers to inject failures in a production environment, like doing Chaos Engineering experiments. Appranix created production-twins allows users to a lot of experiments confidently in a completely different region without affection production setup initially.
	Planned for component level faults to minimize application downtime.Have you validated the application can operate with reduced functionality or degraded performance in the presence of a component failure.	Yes	Same as above, even the component level failures are easier to inject and test with production-twins as opposed to production environments directly. Once customers are comfortable with most of the fault injections using Appranix created production-twins, they can introduce the same failure modes in production environments.

	Planned for dependency failures to minimize application downtime. Have you validated the application can operate effectively in the absence of its dependencies?	Yes	Same as above, it is much easier to validate this Well-architecture with production-twins.
How have you ensured required capacity and services are available in targeted regions? (Azure services and capacity can vary by region, so it is important to understand if targeted regions offer required capabilities.)	Built a capacity model for the application. Is there a capacity model for the application identifying relationships between the utilization of various components, to capture when and how to scale-out?	N/A	N/A
	Planned for expected usage patterns. Has sufficient capacity been provisioned to handle expected usage patterns?	N/A	N/A
	Confirmed Azure service availability in required regions. Are Azure services available in all of the required regions, including any disaster recovery regions?	Yes	Appranix automatically retains the auto-scaling requirements from the production region to target region when a failover happens. Even if the target region does not have the exact resources, for example, the same size VMs in the target region, Appranix will automatically identify the next best VMs and other resources, and adjust the infrastructure-code appropriately for close to guarantee recovery.
	Confirmed Availability Zones are available in required regions. Are Azure Availability Zones available in all of the required regions, including any disaster recovery regions?	No	Appranix automatically adjusts the number of availability zones along with all the load balancer and auto scale sets to get the application running in the target region close to the same performance as that of the production region.
	Validated required capacity is within Azure service scale limits and quotas. Is the required capacity, both initial and future growth, within Azure service scale limits and quotas?	Partially	Appranix helps with simulating the capacity assessments with quick recovery tests. As Azure capacity, and feature capabilities change over time, Appranix helps with automating the capacity assessments with simulated tests.
	Validated all APIs/SDKs against target runtimes and languages for required functionality. Are all APIs/SDKs validated against target runtime/languages for required functionality?	N/A	N/A
	Aligned with Azure roadmaps for required preview services and capabilities. Are any preview services/capabilities required in production?	N/A	N/A

<p>How are you handling disaster recovery for this workload? (Disaster recovery is the process of restoring application functionality in the wake of a catastrophic failure. It might be acceptable for some applications to be unavailable or partially available with reduced functionality for a period of time, while other applications may not be able to tolerate reduced functionality.)</p>	<p>Application is available across multiple regions in an active-active configuration.Is the application and/or its key workloads deployed across multiple regions in an active-active configuration?</p>	Yes	<p>Appranix in these situations only offers backup capabilities of critical application components.</p>
	<p>Application is deployed across multiple regions in an active-passive configuration in alignment with recovery targets.Is the application and/or its key workloads deployed across multiple regions in an active-passive configuration?</p>	Yes	<p>Most of the time, Azure customers may not want to waste their budget on idle cloud resources with active-passive configuration for all the resources. For example, they might be ok with databases with active-passive configuration but not necessarily VMs, containers, load balancers, scale sets, etc. Appranix helps balance the budget with resilience target requirements with active passive and quick failover with active components of the applications. For example, Appranix will manage active-passive databases with appropriate backup functionality and create the rest of the Azure resources on-demand when the DR is required.</p>
	<p>Traffic is routable to the application in the case of a regional failure.</p>	Yes	<p>This is a failure mode analysis. Appranix can do quick test failover and allow customer to check if the traffic is re-routable to another region.</p>
	<p>Defined a backup strategy in alignment with recovery targets.Is application state backed-up and restorable within target recovery times?</p>	Yes	<p>Appranix backups not only data but also the state of the applications, and meta-data so users can achieve complete environment recovery to achieve better recovery targets.</p>
	<p>Defined a disaster recovery strategy to capture recovery steps for failover and failback.Is there a BCDR strategy for the application and/or its key workloads? and are operational runbooks defined for key failure scenarios?</p>	Yes	<p>Appranix automates the entire DR plan creation with Azure native infrastructure-as-code. It then continuously keeps the DR plan refreshed based on the production environment changes. There is no need for specific runbooks or manual DR plans as Appranix continuously reverse engineers the production environments automatically.</p>
	<p>Failover and failback steps and processes are automated. Are operational procedures to failover and failback automated?</p>	Yes	<p>Yes, Appranix takes care of these activities completely.</p>
	<p>Successfully tested and validated the failover and failback approach at least once.Are disaster recovery strategies and operational runbooks for key failure scenarios tested on a regular basis?</p>	Yes	<p>Appranix delivers completely automated DR tests as frequently as the customer want.</p>
	<p>Decomposed the application into distinct subsystems with independent disaster recovery strategies.Has the application been decomposed according to key workloads with their own independent recovery targets and disaster recovery strategies?</p>	Yes	<p>Appranix helps users to decompose large systems into subsystems as separate Cloud Assemblies so users can have independent RPO and RTO targets.</p>

	<p>Network connectivity redundancy for on premise data/application sources. For cross-premises connectivity (ExpressRoute or VPN) are there redundant connections from different locations? i.e. at least two connections from two different locations</p>	<p>Yes</p>	<p>Appranix makes this process much easier than what's possible in the market today as all the recoveries are completely isolated from the production environments in a separate vNet at the target region. Users can also customize the recoveries with post-webhooks to make sure that the on-premise connectivity is re-established appropriately upon a real disaster recovery.</p>
<p>What decisions have been taken to ensure the application platform meets your reliability requirements? (Designing application platform resiliency and availability is critical to ensuring overall application reliability.)</p>	<p>Application processes are stateless. Are all processes stateless with application state externalised?</p>	<p>Partially</p>	<p>Assuming that users have converted all the application components stateless and externalized the data components using Azure services, Appranix, protects stateless components such as containers as well as external data services using appropriate Cloud Assemblies. Changed components are automatically protected using tags. Appranix Application Environment Time Machine captures the past architecture as well as the changed architecture.</p>
	<p>Session state is non-sticky and externalised to a data store. Is session state (if any) non-sticky and externalized to a data store?</p>	<p>N/A</p>	<p>N/A</p>
	<p>Application configuration is treated as code and deployed with the application. Is application configuration treated as code and deployed with the application? i.e. is configuration information captured for operational transparency.</p>	<p>Yes</p>	<p>Appranix captures the point in time metadata or the configurations of the systems. Using this information, the point in time recovery can be performed for all the supported resources in the cloud assembly.</p>
	<p>Application platform services are running in a highly available configuration/SKU. Are all application platform services running in a HA configuration/SKU? e.g. Service Bus Premium.</p>	<p>Yes</p>	<p>Appranix efficiently identifies the respective SKU's and HA's in the recovery region or environment while recovering. If the respective SKU's are not available, then the closest matching SKU is preferred based on the CPU's count.</p>
	<p>Application platform components are deployed across Availability Zones or Availability Sets. Is the application deployed across AZs or ASs?</p>	<p>Yes</p>	<p>Assuming that Azure users have appropriately designed their applications to run on multiple Availability zones, Appranix provides further protection and recovery capabilities across to another region with same level of Availability Zone replication and deployments. For example, if users have deployed their applications across 3 zones in a region when Appranix fails over their application to another region, all their Availability Zone configurations will be retained and Appranix automatically distributes that applications across as many zones as the primary region. In cases where the number of Availability Zones are less in certain regions, Appranix automatically reconfigures to only the available number of zones to make sure the failed over applications can run optimally until the primary/production region to ready for failback.</p>

	<p>Leveraged platform services are Availability Zone aware. Is the underlying application platform service Availability Zone aware? Platform services that can leverage Availability Zones are deployed in either a zonal manner within a particular zone, and/or in a zone-redundant configuration across multiple zones</p>	<p>Yes</p>	<p>Appranix protects the resources either based on the selection or based on the tags. All the resources both the resources in the particular zone and or in zone-redundant configuration across multiple zones are recovered with the zone redundancy configuration preserved during recovery.</p>
	<p>Application platform components are deployed across multiple active regions. Is the application platform deployed across multiple regions? The ability to respond to disaster scenarios for overall compute platform availability and application resiliency is dependant on the use of multiple regions or other deployment locations</p>	<p>Yes</p>	<p>Most of the time, Azure customers may not want to waste their budget on idle cloud resources with active-passive configuration for all the resources. For example, they might be ok with databases with active-passive configuration but not necessarily VMs, containers, load balancers, scale sets, etc. Appranix helps balance the budget with resilience target requirements with active passive and quick failover with active components of the applications. For example, Appranix will manage active-passive databases with appropriate backup functionality and create the rest of the Azure resources on-demand when the DR is required.</p>
	<p>Load balancing is implemented to distribute traffic across multiple nodes. Are application components hosted across multiple nodes based on expected traffic patterns?</p>	<p>Yes</p>	<p>Appranix will recover based on the given configuration in the target region.</p>
	<p>Health probes are implemented to check the health of application components and compound application health. Are health probes used for load balancers to check next hop component health and compound application health for key workloads?</p>	<p>N/A</p>	<p>Appranix will recover the environment with health probes configured by the Azure users.</p>
	<p>Queuing and reliable messaging patterns are used to integrate application tiers. Are queuing and messaging patterns used to asynchronously integrate application components?</p>	<p>N/A</p>	<p>NA</p>
	<p>Client traffic can be routed to the application in the case of region/zone/network outages. How is the client traffic routed to the application in the case of region outage? i.e. is traffic routed through a global load balancer such as Azure Front Door, Azure Traffic Manager or a third-party CDN.</p>	<p>N/A</p>	<p>Appranix recovers the respected resources to the target or recovery regions. Global load balancers can be reconfigured to the newly created setup with a minimal RPO.</p>

	Procedures to scale out application platform components are automated.Are operational procedures to scale out application platform components automated based on key utilization metrics?	N/A	N/A
What decisions have been taken to ensure the data platform meets your reliability requirements? (Designing data platform resiliency and availability is critical to ensuring overall application reliability.)	Data types are categorized by data consistency requirements.Has CAP theorem been applied to the application and key data scenarios to ensure data is categorized according to consistency requirements?	Yes	Appranix performs the backup of the Database VM's configuration and will also do point-in-time data backups. While recovering we would recover the point in time configuration along with the point in time data.
	Data platform services are running in a highly available configuration/SKU.Are data store(s) running in a HA configuration/SKU? i.e. SQL DB in a Zone Redundant configuration, or SQL MI Always-On.	Yes	Appranix automatically captures zone redundant services and recover them in another region with the same configuration.
	Data is replicated across multiple regions.Is data replicated across paired regions?	Yes	Appranix manages the data replication across the region using Azure native capabilities.
	Data is replicated across Availability Zones.Is data replicated across Availability Zones within a region?	Yes	Appranix enables Azure customers implement data replication between regions and manages them for recovery
	Data is backed-up on zone/geo-redundant storage.Is data backed-up on zone/geo-redundant storage?	Yes	Appranix natively covers this capability for backup, recovery and DR
	Active geo-replication is used for data platform components such as storage and databases.Is active geo-replication used?	Yes	Appranix natively covers this capability for backup, recovery and DR
	Application traffic can be routed to data stores in the case of region/zone/network outages.How is application traffic routed to data sources in the case of region/zone/network outage? i.e. Cosmos DB Automatic Failover.	N/A	N/A

	Read operations are segregated from update operations. Have read operations been segregated from update operations across application data stores?	N/A	N/A
	Load balancer health probes assess data platform components. Do health probes assess critical internal data dependencies? e.g. do health probes perform a mutable database operation.	N/A	N/A
	Data restore processes have been defined to ensure consistent application state when data is corrupted or deleted. Has a data restore process been defined to ensure a consistent state for the application?	Yes	Appranix natively covers this capability for backup, recovery and DR
	Data restore processes have been validated and tested to ensure consistent application state when data is corrupted or deleted. Has a data restore process been validated and tested to ensure a consistent state for the application?	Yes	Appranix natively covers this capability for backup, recovery and DR
How does your application logic handle exceptions and errors? (Resilient applications should be able to automatically recover from errors by leveraging modern cloud application code patterns.)	Have a method to handle faults that might take a variable amount of time to recover from. There can also be situations where faults are due to unanticipated events, and that might take much longer to fix. These faults can range in severity from a partial loss of connectivity to the complete failure of a service.	N/A	N/A
	Request timeouts are configured to manage inter-component calls. Have you configured timeouts for inter-component calls?	N/A	N/A

	<p>Retry logic is implemented to handle transient failures, with appropriate back-off strategies to avoid cascading failures. Have you implemented retry logic to handle transient application failures as well as transient failures with internal or external dependencies?</p>	N/A	N/A
	<p>The application is instrumented with semantic logs and metrics. Is the application instrumented with semantic logs to capture, alert and respond to errors?</p>	N/A	N/A
<p>What decisions have been taken to ensure networking and connectivity meets your reliability requirements? (Identifying and mitigating potential network bottlenecks or points-of-failure supports a reliable and scalable foundation over which resilient application components can communicate.)</p>	<p>All single points of failure have been eliminated from application communication flows. Have all single points of failure been eliminated from the data path (on-premises and Azure)? i.e. no single VM NVA or no single ExpressRoute connection</p>	Yes	<p>Appranix would recover point in time configuration of the network. Any misconfigurations can be fixed in the target by recovering the network with the correct configuration.</p>
	<p>Health probes are configured for Azure Load Balancer(s) to assess application traffic flows and compound health. Are Azure Load Balancer health probes configured to assess traffic flows and compound application health?</p>	Yes	<p>Appranix protects the loadbalacner configuration including the health configuration, backend pool configuration and recover it in the DR region with the same configuration as is in the production environment.</p>
	<p>Azure Load Balancer Standard or Zone redundant application gateways are used to load balance traffic across Availability Zones. Are there any mitigation plans defined in case data size exceeds limits? i.e. purging or archiving</p>	Yes	<p>Appranix protects both standard and zone redundant application gateways and can recover it either in the same region or in the cross region based on the usecase along with the backend pools associated with the Loadbalancer and the Application gateway.</p>
	<p>Redundant connections from different locations are used for cross-premises connectivity (ExpressRoute or VPN). Are there redundant connections from different locations for cross-premises connectivity (ExpressRoute or VPN)? i.e. at least two connections from two different locations.</p>	N/A	NA

	A failure path has been simulated for cross-premises connectivity.Has a failure path been simulated to ensure connectivity is available over alternative paths? i.e. using S2S as a backup for Express Route.	N/A	NA
	Zone redundant gateways are used for cross-premises connectivity (ExpressRoute or VPN).Are ExpressRoute/VPN zone redundant gateways being used?	N/A	NA
	Network traffic is monitored, and a response plan is in place to address network outages.Is network traffic and connectivity monitored? and are operational procedures in place to respond to network outages?	Yes	Having a production twin created using Appranix, you could predetermine the failure scenairos by running network load test and have the production appropriately configured for scaling.
What reliability allowances for scalability and performance have you made? (Resilient applications should be able to automatically scale in response to changing load to maintain application availability and meet performance requirements.)	The application has dedicated cross-premises bandwidth.Does the application have dedicated cross-premises bandwidth? Or is bandwidth shared with other applications?	N/A	N/A
	Components with sensitive latency requirements are collocated.Are there any components or scenarios that are very sensitive to latency and require resources to be collocated in close proximity?	N/A	N/A
	Gateways (ExpressRoute or VPN) have been sized according to expected cross-premises network throughput.Have gateways (ExpressRoute or VPN) been sized accordingly to the expected cross-premises network throughput?	N/A	N/A

Expected throughput passing through security/network appliances has been tested and autoscaling is configured based on throughput requirements.Is autoscaling enabled and integrated within Azure Monitor?	Partially	Appranix lets you run security scanning and testing in the 'production twin' that is exactly same as production but in a different region so that all the vulnerability scanning can be done without disturbing the production but with the production data and configuration.
Autoscaling is enabled for application components and integrated with Azure Monitor.Is autoscaling enabled and integrated within Azure Monitor?	Partially	Appranix automatically protects those instances that are created as part of autoscaling without any manual intervention. Appranix application environment time machine exactly knows how many instances were running with what configuration at a particular point in time.
Autoscaling has been tested and the time to scale in/out has been measured.Has autoscaling been tested? and has the time to scale in/out been measured?	N/A	N/A
Tested and validated defined latency and defined throughput targets per scenario and component.Are latency and throughput targets defined, tested and validated per scenario/service? i.e. first byte in to last byte out	N/A	N/A
Calculated target data sizes and associated growth rates per scenario and component.Are target data sizes and associated growth rates calculated per scenario/services?	N/A	N/A
Operational procedures are defined in case data sizes exceed limits.Are there any mitigation plans defined in case data size exceeds limits? i.e. purging or archiving	Partially	Appranix ability to create a production twin helps to do stress and load testing to predetermine the capacity and plan properly to be able to handle the load in real-time.
Validated that long-running TCP connections are not required for the workload.Does the workload require a large amount of long running TCP connections which may cause SNAT port exhaustion?	N/A	N/A

	Throttling is implemented to govern inbound application calls and inter-component calls. Have you implemented throttling for all inbound and inter-component calls?	Partially	Appranix ability to create a production twin helps to do stress and load testing and helps in predetermine the capacity and plan properly to be able to handle load both the data and the in bound requests.
What reliability allowances for security have you made? (Identifying and addressing security-related risks helps to minimize application downtime and data loss caused by unexpected security exposures.)	The identity provider (AAD/ADFS/AD/Other) is highly available and aligns with application availability and recovery targets. Is the identity provider (AAD/ADFS/AD/Other) and network connectivity to the identity provider highly available?	Yes	Appranix allows protecting the identity provider (AAD/ADFS/AD) and recover the entire active directory forest configuration and allows the DR environment to have an Active Directory setup seamlessly similar to the production environment.
	All external application endpoints are secured? i.e. Firewall, WAF, DDoS Protection Standard Plan, etc. Are all external application endpoints secured using security services? Such as Azure Firewall, Application Gateway and Azure Front Door WAF, or DDoS Protection Standard Plan. How threat vectors such as DDoS attacks are mitigated ultimately has a bearing on application reliability	Yes	Appranix supports configuring the post recovery webhook using which all automation around connecting to the external services can be taken care of. Appranix provides protection and recovery capabilities using the cloud-native approach and hence all the firewall rules to the external endpoints etc in the DR region will be exactly the same as in the production environment at that point in time.
	Communication to Azure PaaS services secured using Virtual Network Service Endpoints or Private Link. Is communication to Azure PaaS services secured by using VNet Service Endpoints or Private Link?	Yes	As mentioned above, Appranix recovery can be configured with the post recovery webhook to connect the external script that auto-configures all the external services.
	Keys and secrets are backed-up to geo-redundant storage. Are Keys and secrets backed-up to geo-redundant storage?	Yes	Appranix supports the protection of Azure vault and allows to store both the data and the metadata into the geo-redundant storage account.
	The process for key rotation is automated and tested. Is the process for key rotation is automated and tested?	N/A	N/A

	Emergency access break glass accounts have been tested and secured for recovering from Identity provider failure scenarios.Are break glass AAD accounts setup and secured?	Partially	Identity providers can be protected and recovered in a complete working state in a different region within the same Azure subscription.
What reliability allowances for operations have you made? (Operational processes for application deployment, such as roll-forward and roll-back, should be defined, sufficiently automated, and tested to help ensure alignment with reliability targets.)	Application can be automatically deployed to a new region without any manual operations to recover from disaster scenarios.Can the application be deployed automatically from scratch without any manual operations?	Yes	Appranix allows to configure the CI/CD pipeline and hence the creation of the environment in the recovery region can be completely automated.
	Application deployments can be rolled-back and rolled-forward through automated deployment pipelines.Can N-1 or N+1 be deployed via automated pipelines where N is current deployment in production?	Yes	Appranix protection can be configured in the CI/CD pipeline to take a snapshot of the complete Application environment along with the application data and the infrastructure metadata before the deployment. This helps in achieving instant rollback of a bad deployment.
	The lifecycle of the application is decoupled from its dependencies.Is the lifecycle of the application is decoupled from its dependencies?	Yes	Appranix allows automatic reset of the recovered Application environment and the same can be configured to the pipeline. This allows reducing the cloud cost of deployment testing by cleaning up the environment after generating the required test results.
	The time it takes to deploy an entire production environment is tested and validated.Is the time it takes to deploy an entire end-to-end production environment understood and validated?	Yes	Appranix enables users to test entire environment recovery very quickly. Moreover, users can configure recovery simulation policies to automate periodic testing. Appranix also provides Resilience Insights where the application recovery time is shown and can be optimized to ensure better RTO.
How do you test the application to ensure it is fault tolerant? (Application workloads should be tested to validate reliability against defined reliability targets.)	The application is tested against critical Non-Functional requirements for performance.Is the application tested for performance?	Yes	Appranix provides support to create production twins for the purposes of test and DR with minimal RTO.
	Load Testing is conducted with expected peak volumes to test scalability and performance under load.Is Load testing performed using expected peak volumes? i.e. do you test application scalability and performance under load?	Yes	Appranix helps to create the production twin which can be stressed and load tested instead of directly testing in the production environment to identify application performance related issues. The production twin can be obtained by recovering the environment in the same or the DR region based on the testing requirement.

	Chaos Testing is performed by injecting faults.Is the system tested under critical failure scenarios to validate availability and recovery targets are fully satisfied?	Yes	Chaos testing can be done on the production twin. Also, the backups can be tested in the recovery simulation or the DR region after recovery.
	Tests are automated and carried out periodically or on-demand.Are tests automated and carried out periodically or on-demand?	Yes	Application test automation can be combined with Appranix recovery simulation. Once the recovery is completed, a post recovery webhook action can trigger the test automation and save the results to a blob store. The reset of the recovered environment is automated so that the entire test suite from the creation of an environment, the performance of the test, and the reset of the environment can be completely automated.
	Critical test environments have 1:1 parity with the production environment.Do key test environments have 1:1 parity with the production environment?	Yes	Appranix creates the 1:1 parity with the production environment with a touch of a button including the VMs, containers, IP address, Firewall rules, Load Balancer configuration, gateways, etc.
	The application is instrumented with semantic logs and metrics.Is the application instrumented with semantic logs and metrics using APM technologies such as Application Insights?	N/A	N/A
	Application logs are correlated across components.Are application logs correlated across components to ensure end-to-end system flows and component dependencies are monitored?	N/A	N/A
	All components are monitored and correlated with application telemetry.Are all components monitored and correlated with application level telemetry?	N/A	N/A
	Key metrics, thresholds, and indicators are defined and captured.Have key metrics, thresholds and indicators been defined and captured?	N/A	N/A

<p>How do you monitor and measure application health? (Monitoring and measuring application availability is vital to qualifying overall application health and progress towards defined reliability targets.)</p>	<p>A health model has been defined based on performance, availability, and recovery targets and is represented through monitoring dashboard and alerts. Has a health model been defined to qualify what 'healthy' and 'unhealthy' states represent across all application components, in a measurable and observable format? and is this health model fully represented through monitoring dashboards and alerts?</p>	Partially	<p>Azure Cloud Application environment enabled with the health model and protected with Appranix is a great approach to ensure Business continuity. The health model dashboard would provide detailed insights into the environment's ability to recover from the failure. The Legacy backup would bring up the data, but spending days configuring the network configurations, firewall rules, load balancer rules and instance size mapping in the recovery region is going to be a lot of work. Appranix addresses it with a click of a button and with no or very minimal pre and post recovery steps, you could make sure that the Azure cloud application environment is up and running in the health dashboard. The difference between the legacy backup tools vs Appranix is the end-to-end recovery time and manual efforts between your Application going down on an event in the health dashboard and the Application coming back up online in the health dashboard after recovering.</p>
	<p>Azure Service Health events are used to alert on applicable Service level events. Have Azure Service Health events been integrated to alert on applicable Service level events?</p>	Partially	<p>Appranix supports end-to-end API integration. In the case of Azure service health events in a region or zone, the API integration to the monitoring systems could automatically trigger a recovery operation either to a different region based on the failure type along with the pre and post-recovery webhooks that would bring back the Application with very minimum RTO.</p>
	<p>Azure Resource Health events are used to alert on resource health events. Are Azure Resource Health event used to alert on resource health events? An appropriate threshold for resource unavailability must be set to minimise signal to noise ratios so that transient faults do not generate an alert. For example, configuring a virtual machine alert with an unavailability threshold of 1 minute before an alert is triggered.</p>	N/A	N/A
	<p>Monitor long-running workflows for failures. Do you monitor any long running application workflows and operationalise failure scenarios?</p>	N/A	N/A