ForeScout

Highlights:

- Provide resiliency for CounterACT services
- Protect against single- or multi-point failures
- Support centralized and distributed CounterACT deployments
- Automate failover and intelligent reallocation of workloads
- Help meet IT service continuity mandates
- Enable cross-site failover for disaster recovery scenarios
- Perform manual failover to facilitate maintenance procedures and upgrades
- Avoid excessive cost and complexity of idle, standby appliances

ForeScout CounterACT® Failover Clustering

Service continuity and resiliency for single- or multi-site deployments

ForeScout CounterACT® is a physical or virtual security solution that is designed to operate as a critical service in your enterprise security environment. CounterACT helps you see the devices that connect to your network, control their access to network resources based on security policy and compliance state, and orchestrate security workflows based on comprehensive, real-time device intelligence. These functions rely on the availability and uptime of CounterACT services. Thus, any extended interruption can compromise your security posture and impact business operations.

As with any critical service, organizations need to consider deployment architectures that are resilient to system failures, site-wide disruptions and natural or human-induced disasters. Planning and implementing a recovery strategy reduces downtime and enables the continuation of vital business and security systems, such as CounterACT. For these and many other availability benefits, ForeScout supports service resilience through failover clustering.

CounterACT Failover Clustering

Most CounterACT deployments involve multiple physical or virtual appliances, sometimes distributed across several sites. Each appliance can provide a range of services—device visibility, posture assessment, access control and policy enforcement—for a number of endpoints. Failover clusters harness the unallocated processing capacity in these appliances to provide service resiliency without the added cost and complexity of idle, standby appliances.

Failover clustering lets you create logical groups of appliances and implements an automated process for reallocating the workload(s) of one or more failed nodes, a cluster or even an entire site. Clusters can provide resiliency for centralized or distributed deployments and can be deployed in single- or multi-site environments. They can help ensure CounterACT service availability to provide the service continuity necessary to help meet IT operational and security mandates.

How Failover Clustering Works

All appliances in a failover cluster must be managed by the same Enterprise Manager. Deployments should be planned in such a way that appliances have extra capacity to receive the anticipated failover workload share (the failover assignment) in addition to their own normal workload (the original assignment).

A failover event is triggered when the Enterprise Manager cannot connect to an appliance for a defined period of time. When an appliance (or a cluster or site) is determined to have failed, its workload is transferred to, and load balanced across, the assigned receiving appliances. Failback occurs after a failed appliance reconnects to the Enterprise Manager and regains the management of endpoints and network devices previously transferred to the recipient appliances.

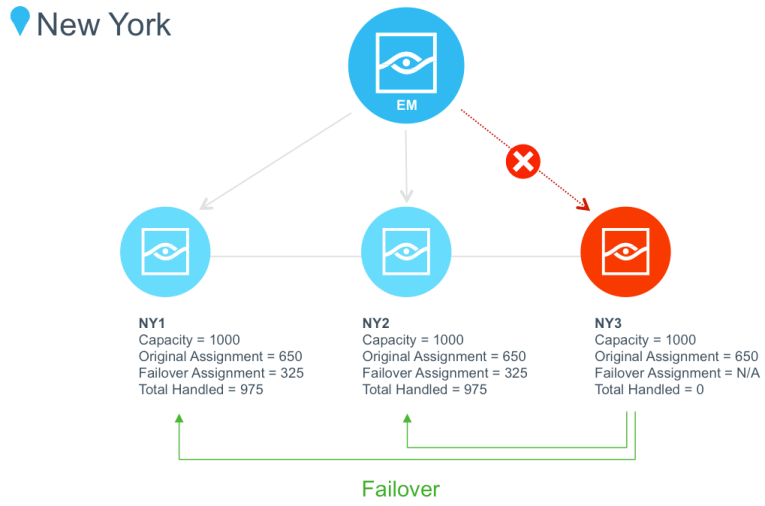

Figure 1: A three-appliance failover cluster for a New York data center, showing the original workload distribution and spare capacity provisioning for a single-node failover.

Routine maintenance procedures and software upgrades on CounterACT appliances do not trigger automated failover. However, you can initiate manual failover of one or more appliances while maintenance work is being performed on them to enable service continuity. Once maintenance procedures are complete, you can manually fail back the appliance(s).

Cross-Cluster and Cross-Site Failover

In addition to failover and workload distribution between appliances within a single cluster, you can also configure failover scope to extend resilience across multiple clusters and locations. When an appliance fails, its workload is first distributed to other nodes within the cluster that have available capacity. When all in-cluster capacity is allocated, workloads are then distributed to appliances in other clusters in the failover scope. This also enables cross-site failovers in the event of an entire cluster or site failure for disaster recovery purposes.

Figure 2 shows a failover scope for our New York cluster that includes recipient appliances in a separate cluster located in New Jersey. If a New York appliance fails, its workload is distributed across the two clusters, with a preference for local receivers. If the entire New York site fails, its assignments are distributed across appliances in New Jersey.

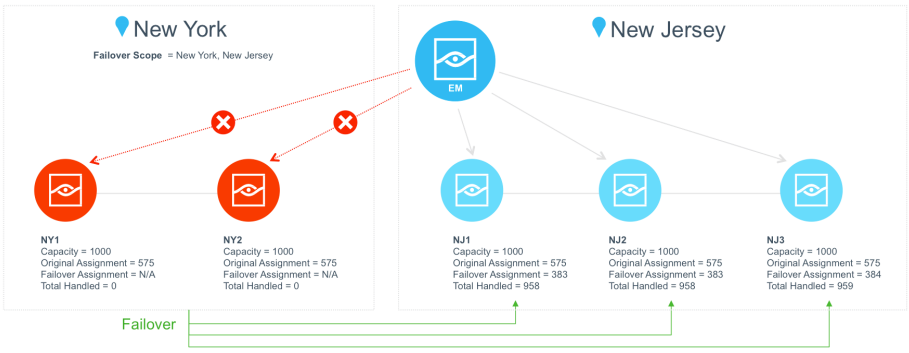

Figure 2: Failover scope for multi-site failover.

Licensing

Failover Clustering is an add-on capability for ForeScout CounterACT that is licensed and sold separately.

Learn more at
www.ForeScout.com

ForeScout

ForeScout Technologies, Inc.
190 West Tasman Drive
San Jose, CA 95134 USA

Toll-Free (US) 1-866-377-8771
Tel (Intl) +1-408-213-3191
Support 1-708-237-6591